Dandelion: Privacy-Preserving Transaction Propagation in Bitcoin's P2P Network

Presenter: Giulia Fanti

Joint work with: Shaileshh Bojja Venkatakrishnan, Surya Bakshi, Brad Denby, Shruti Bhargava, Andrew Miller, Pramod Viswanath

Bitcoin P2P Primer

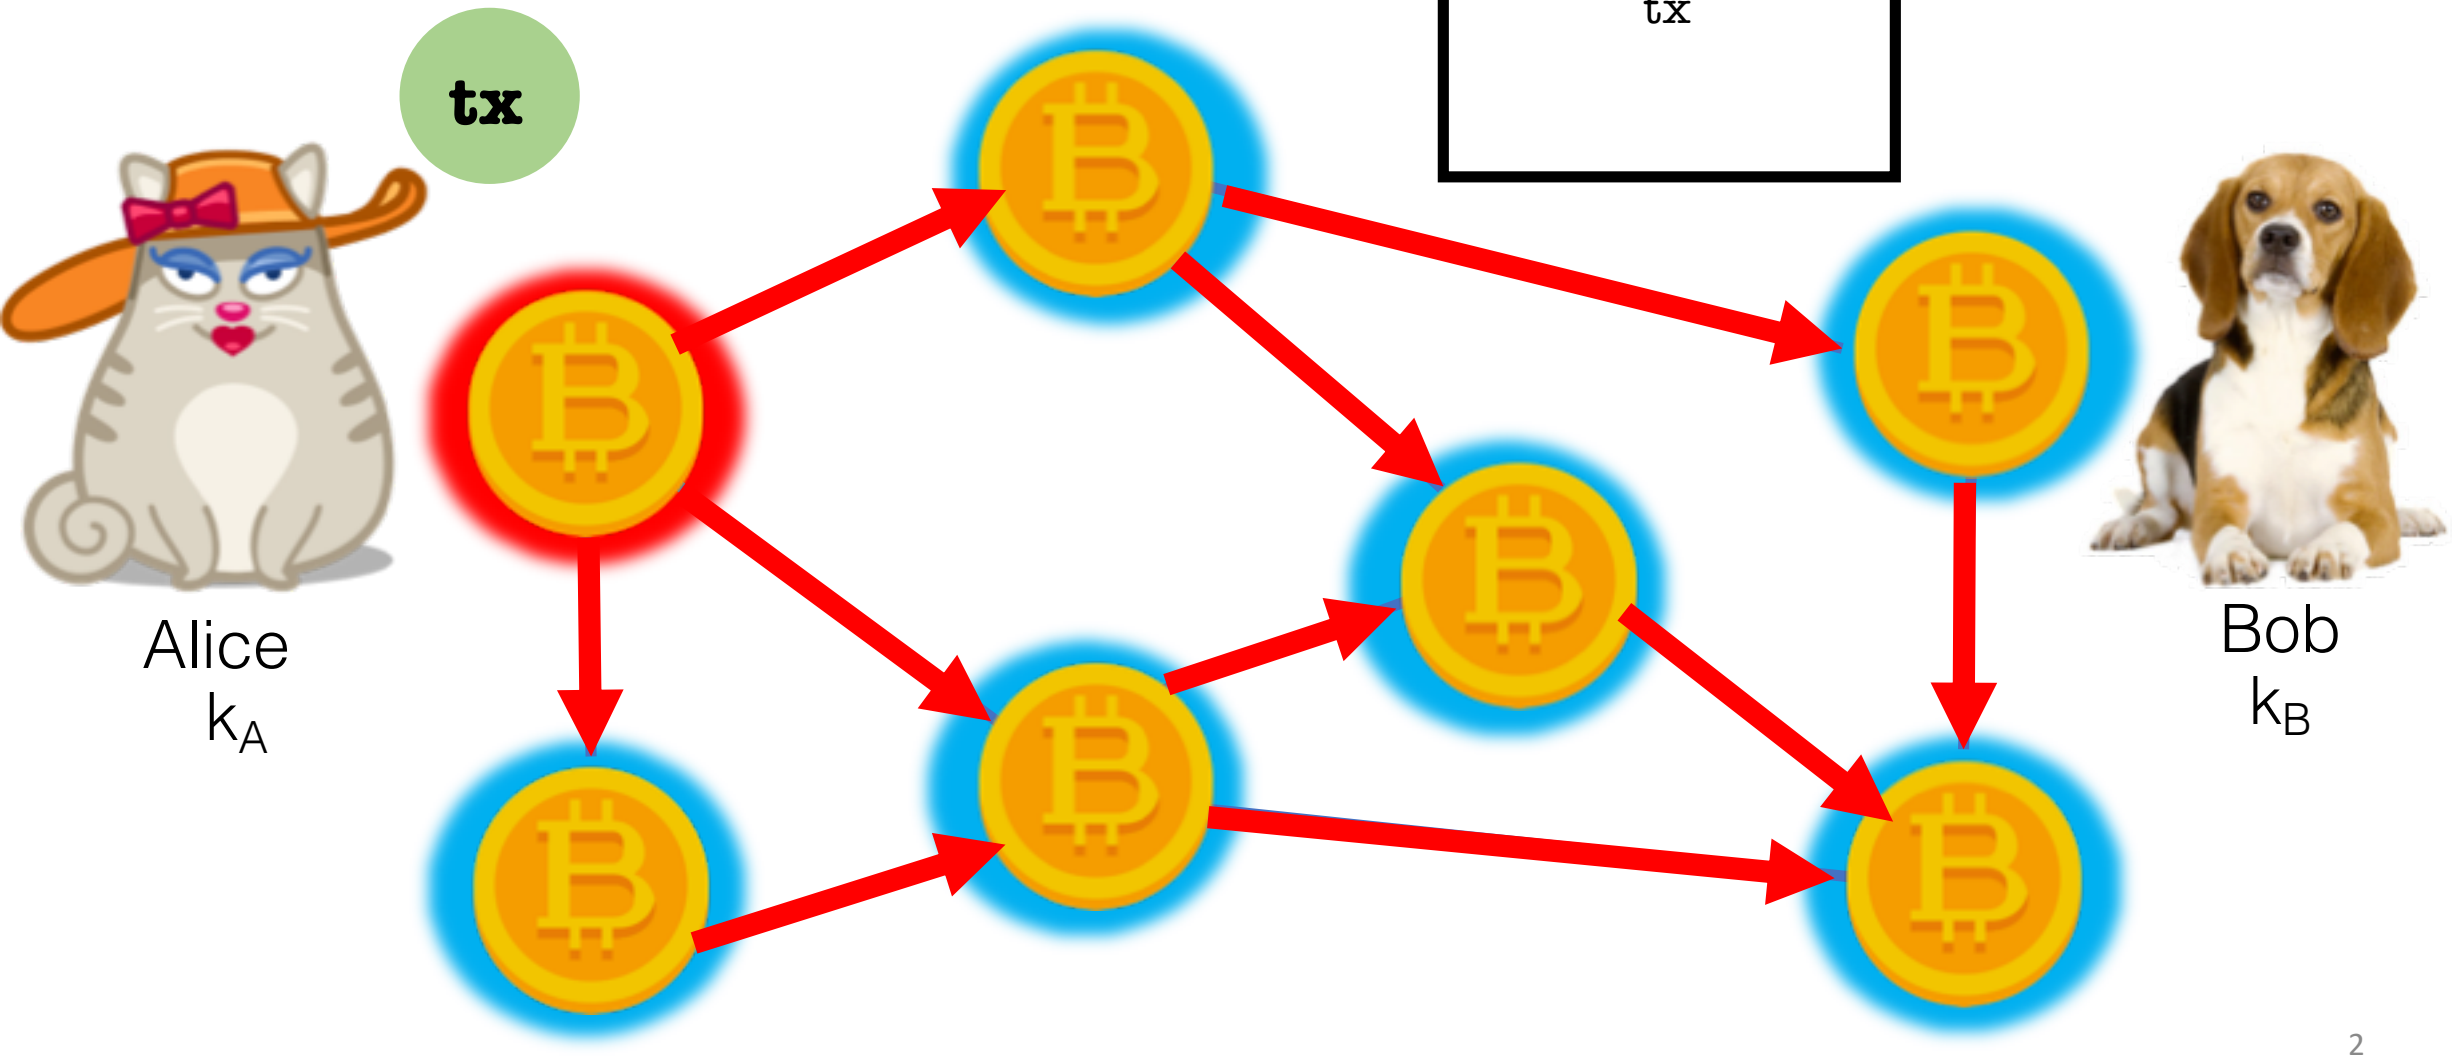

Privacy requirement:

Address and real identity must be unlinkable

Today, messages spread with **diffusion.**

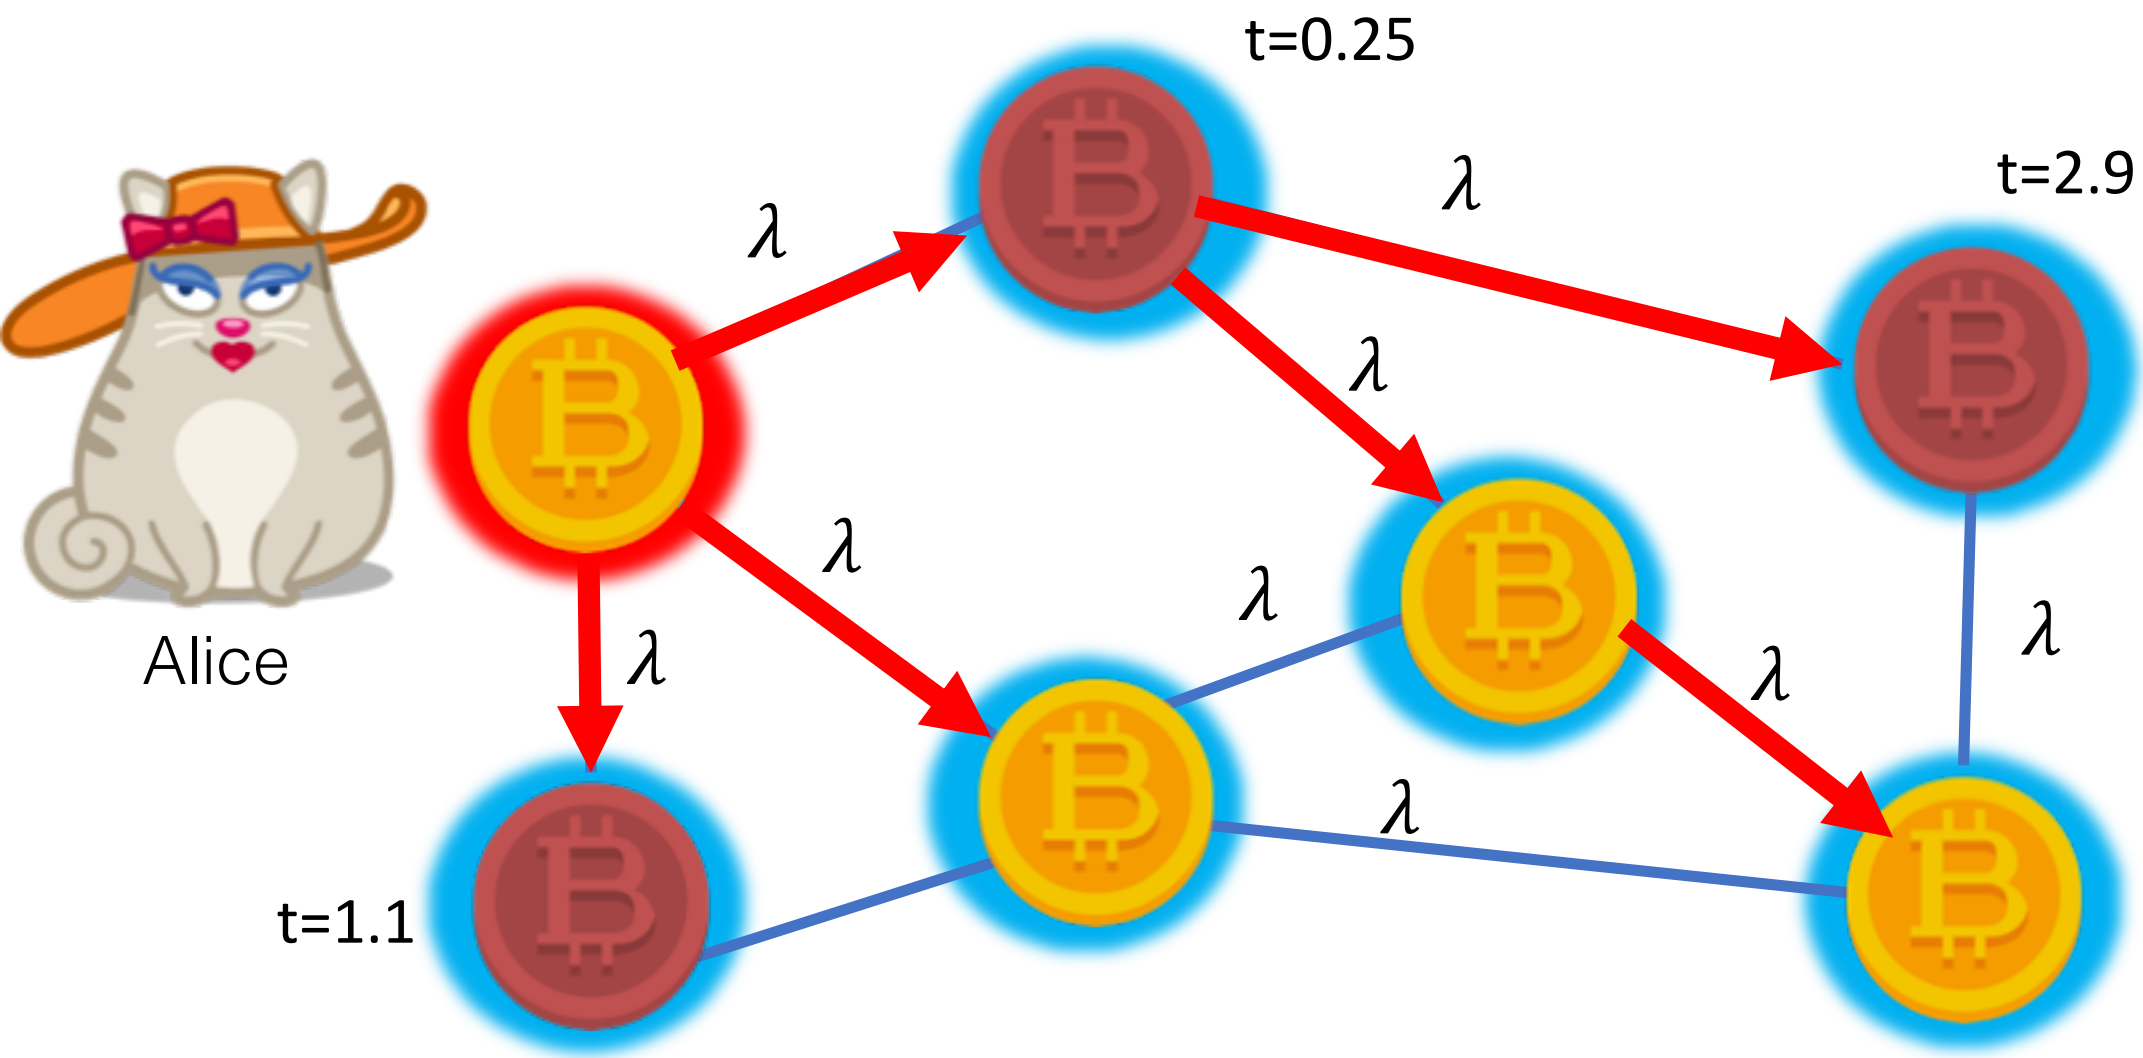

Diffusion is vulnerable to source detection!

Biryukov et al. CCS 2014
Koshy *et al.,* Financial Crypto 2014
F. and Viswanath, NIPS 2017

Dandelion

Lightweight transaction propagation algorithm with provable privacy guarantees.

Venkatakrishan *et al.*, ACM Sigmetrics 2017; F. *et al.*, ACM Sigmetrics 2018

FAQ: Why not alternative solutions?

Connect through Tor

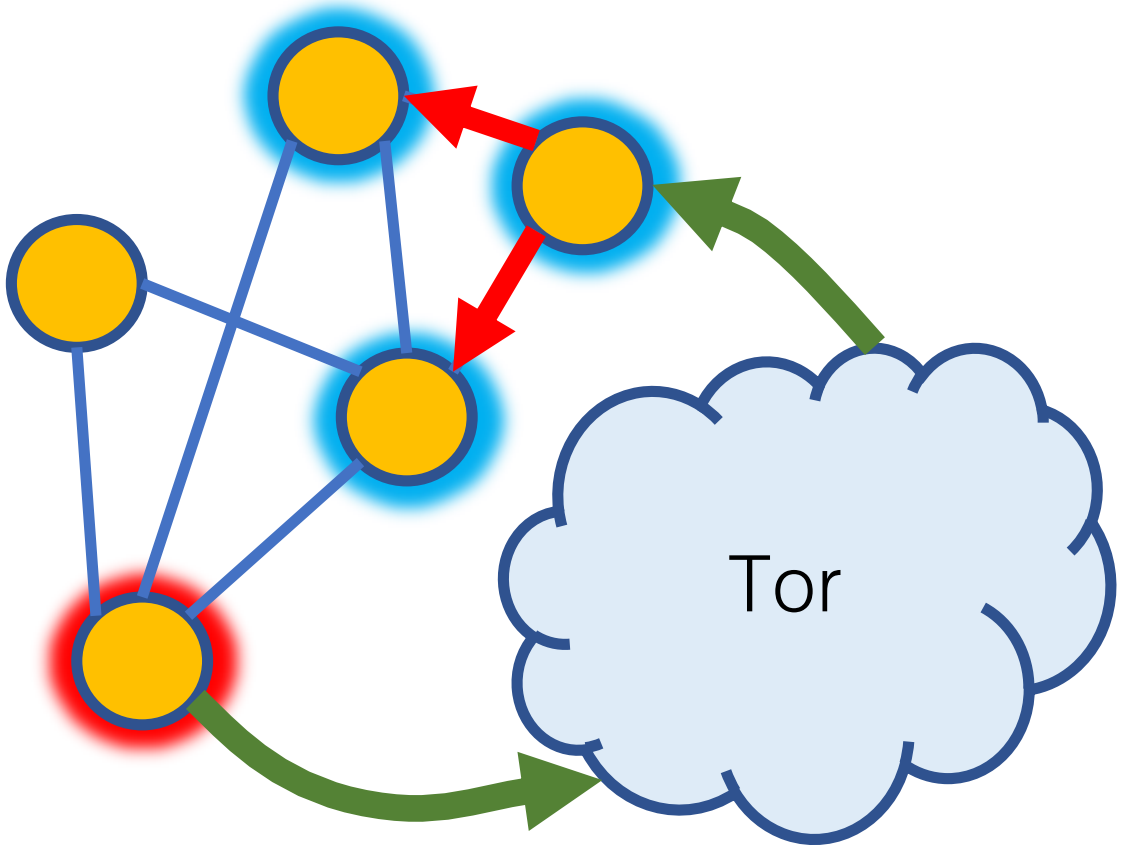

I2P Integration (e.g. Monero)

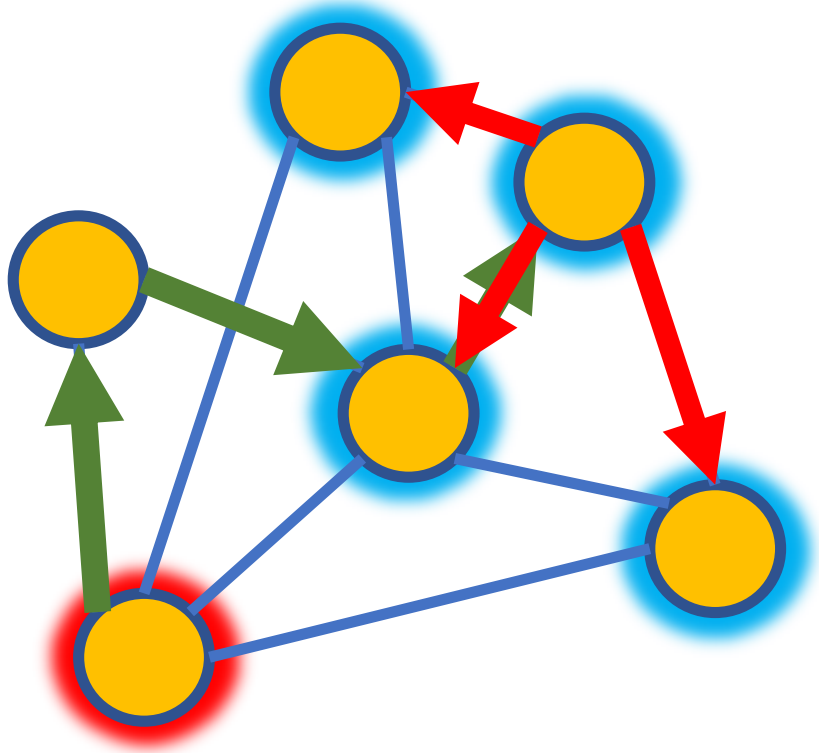

Model

Assumptions and Notation

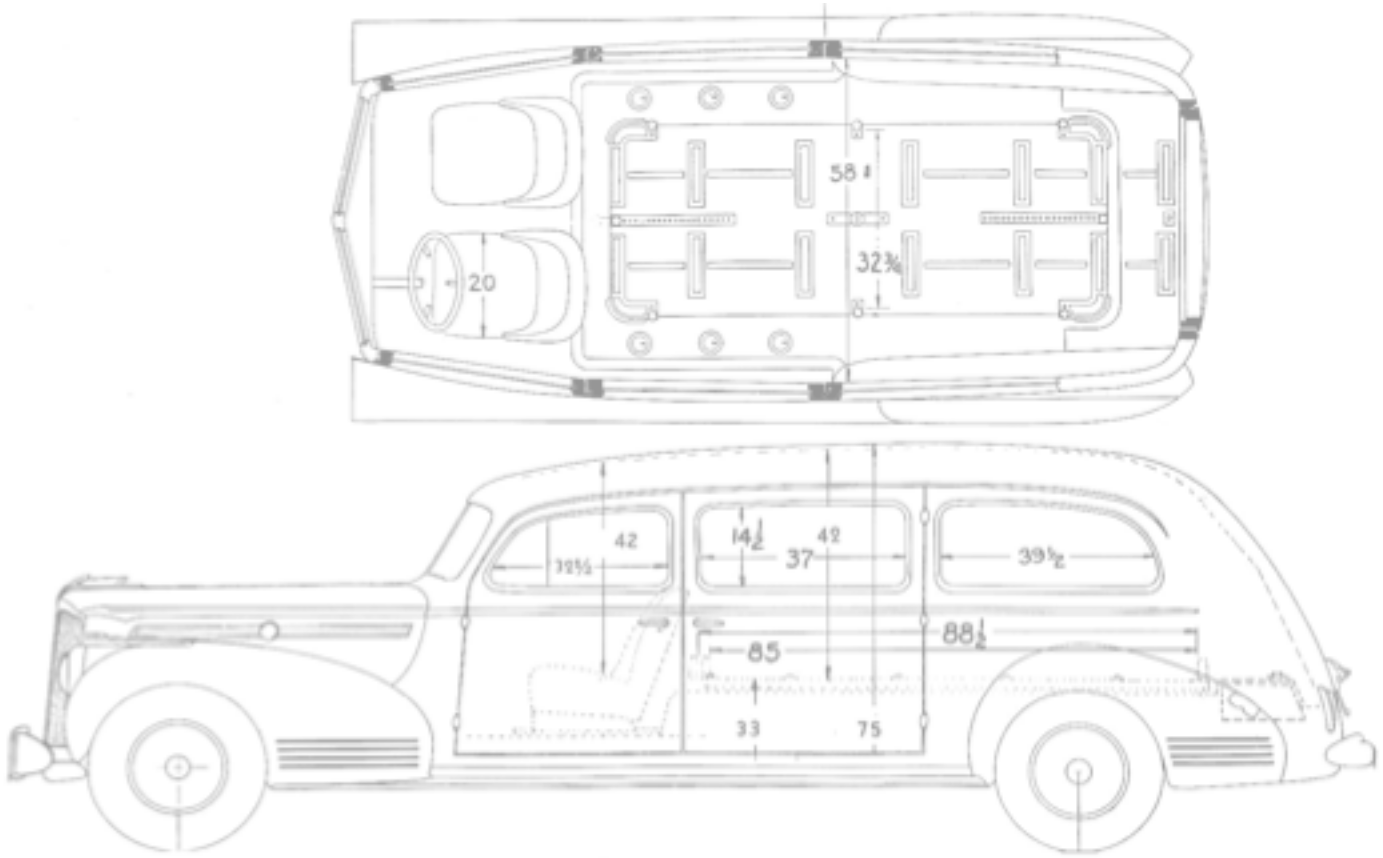

Adversarial model

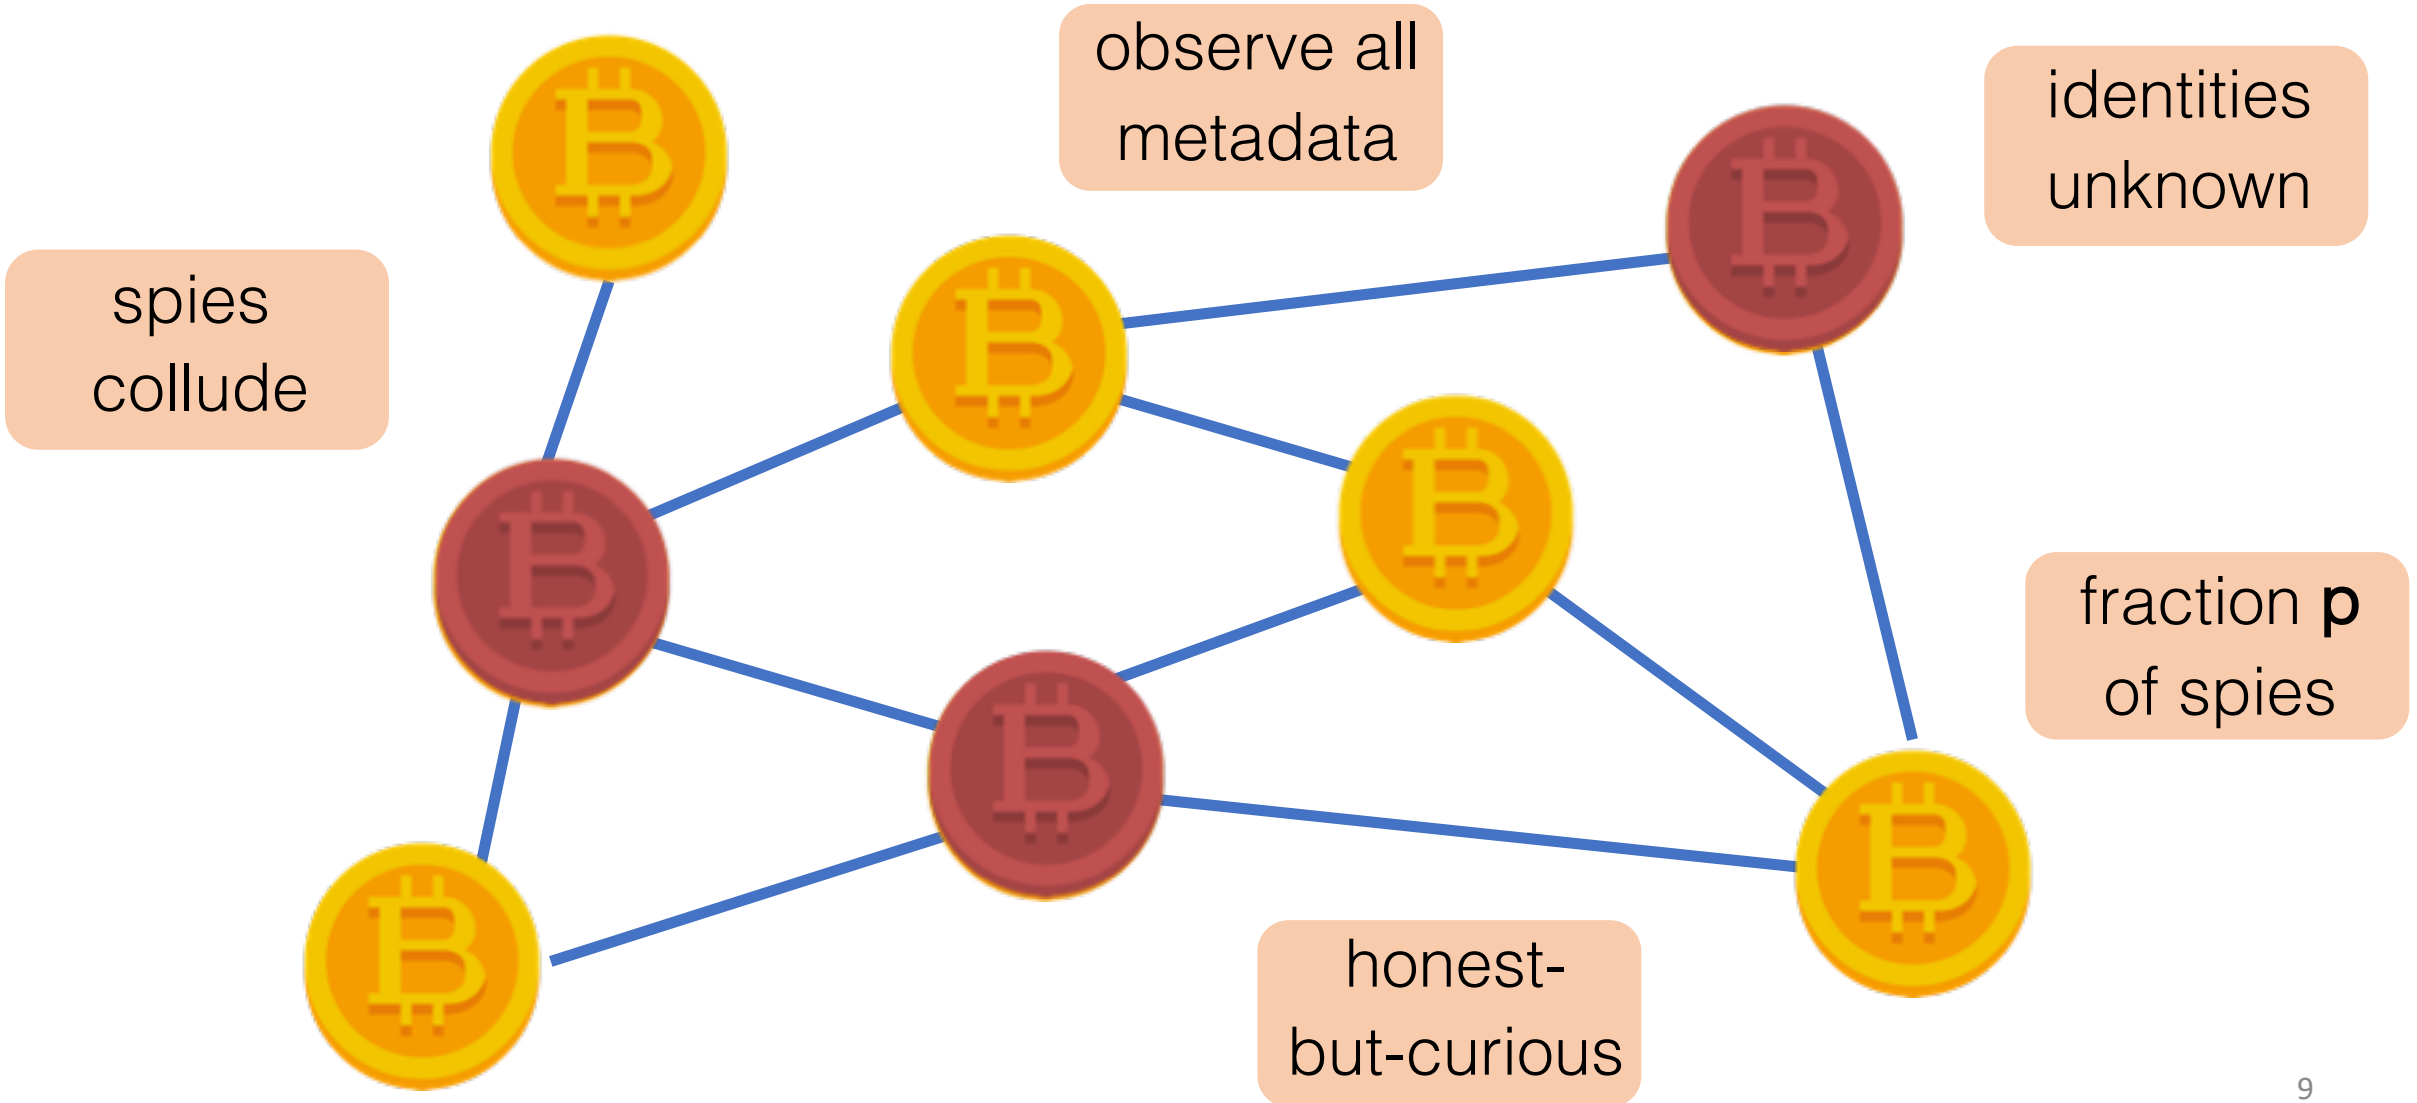

Metric for Anonymity

Recall

$$\frac{1}{n}\sum_{v} 1\{M(v\text{'s tx}) = v\}$$

Number honest users

User

Mapping

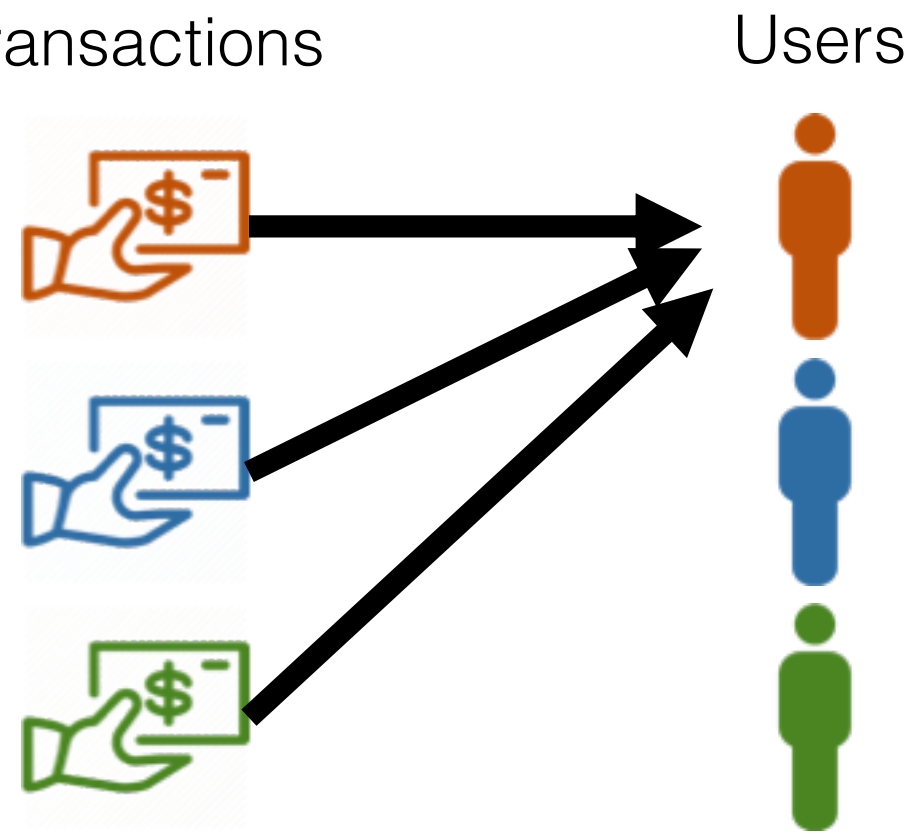

Precision

$$\frac{1}{n}\sum_{v} \frac{1\{M(v\text{'s tx}) = v\}}{\#\text{ tx mapped to v}}$$

Goal:

Design a distributed flooding protocol that minimizes the maximum precision and recall achievable by a computationally-unbounded adversary.

Fundamental Limits

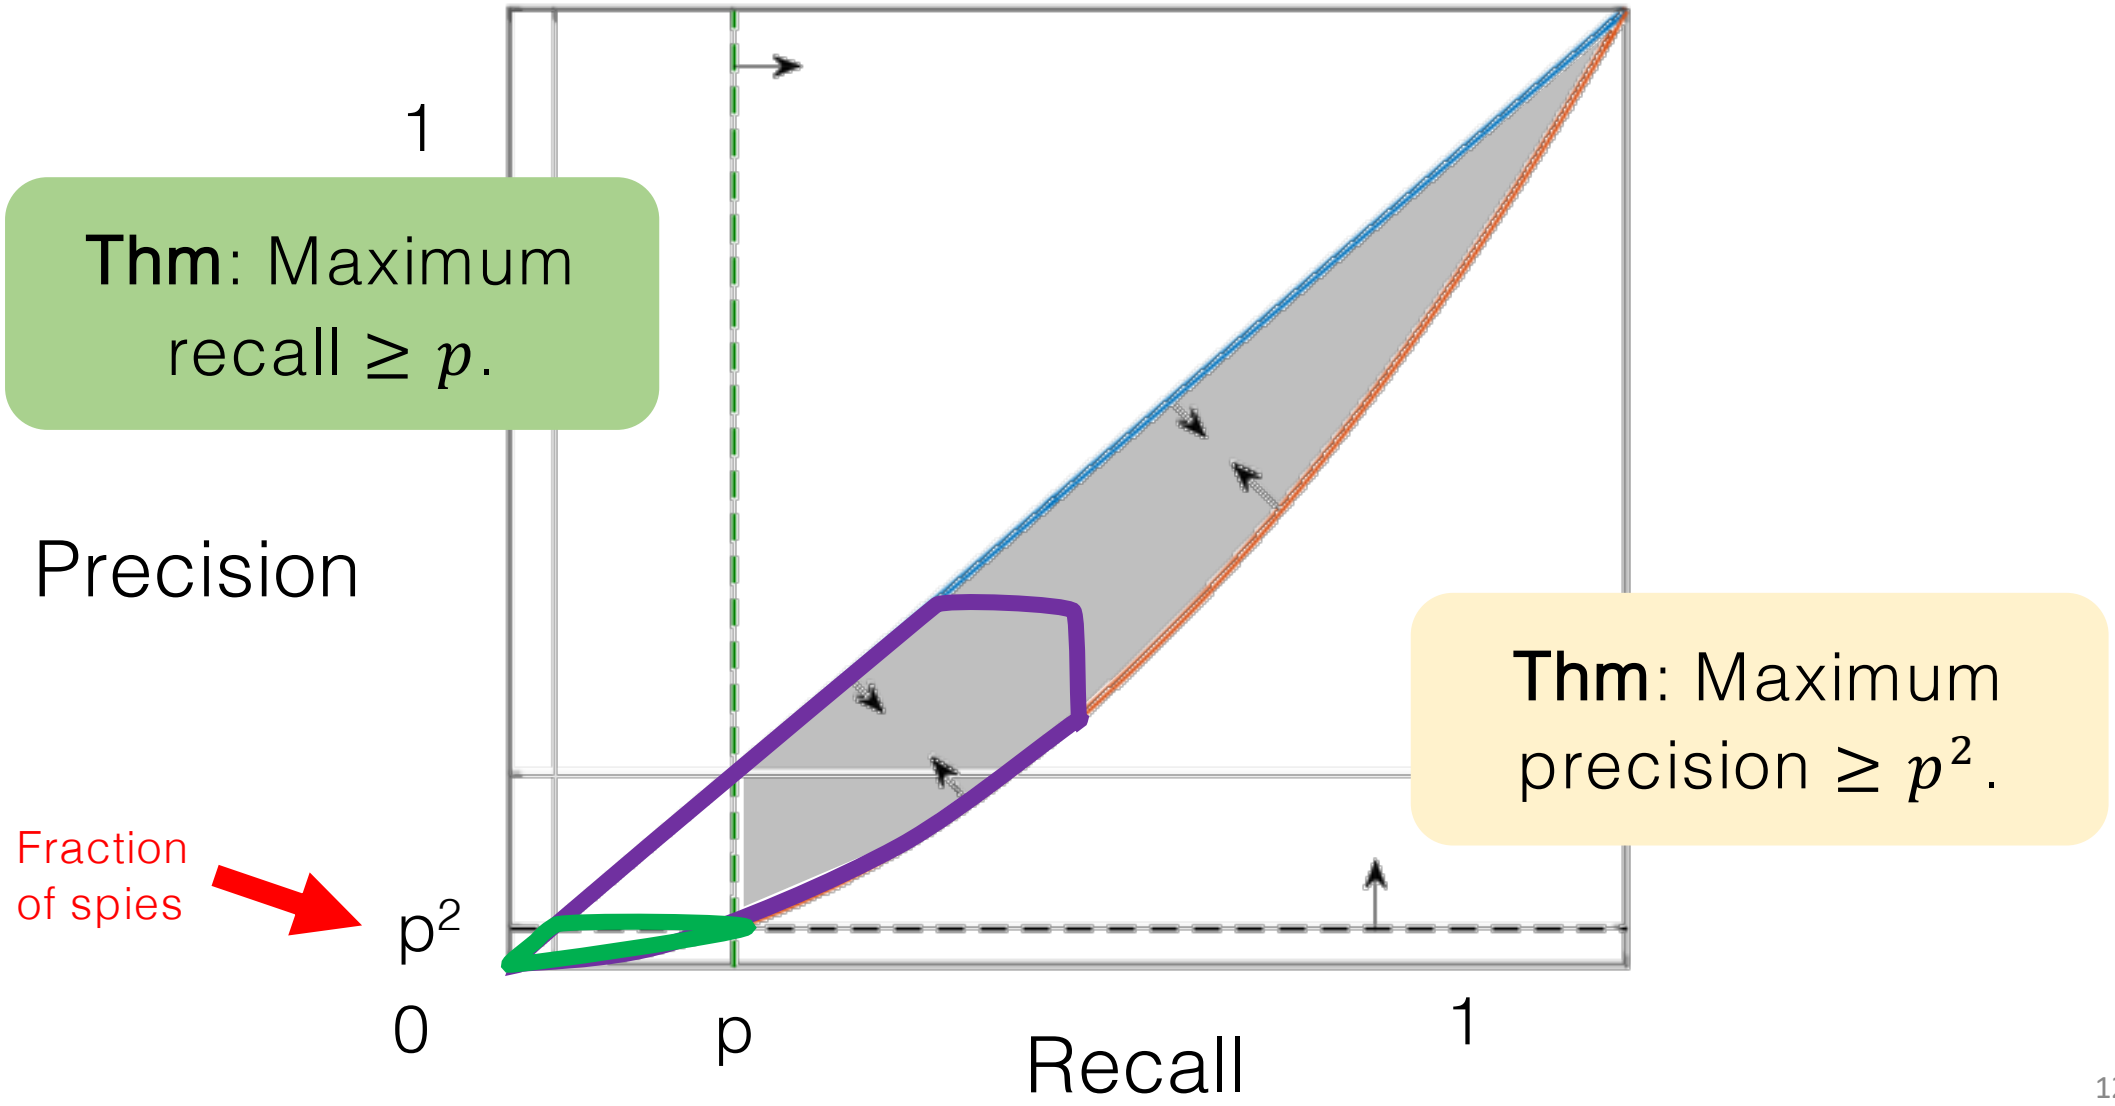

Thm: Maximum recall $\geq p$.

Thm: Maximum precision $\geq p^2$.

What are we looking for?

Asymmetry

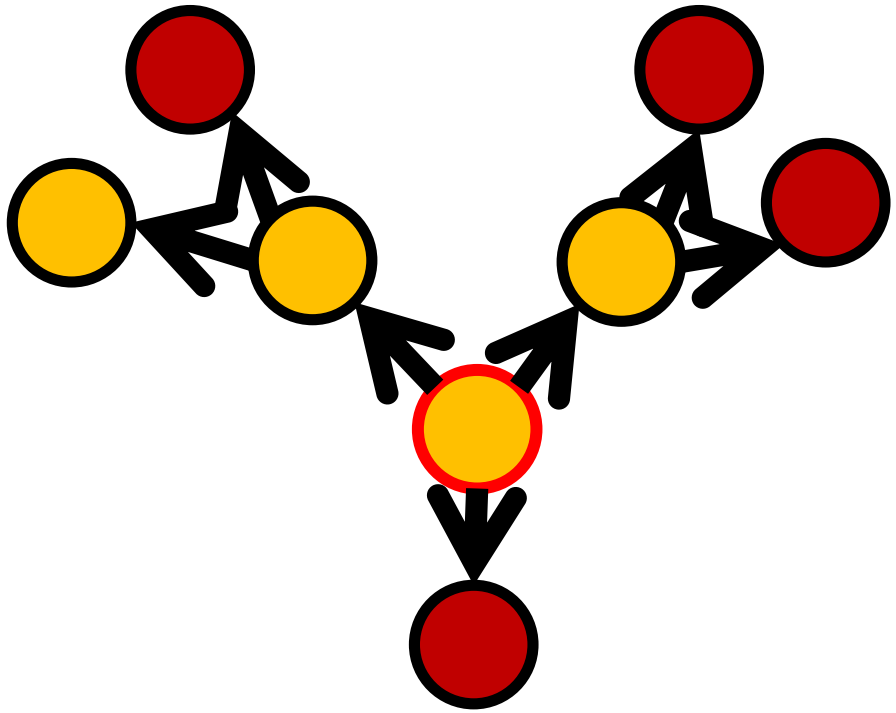

Mixing

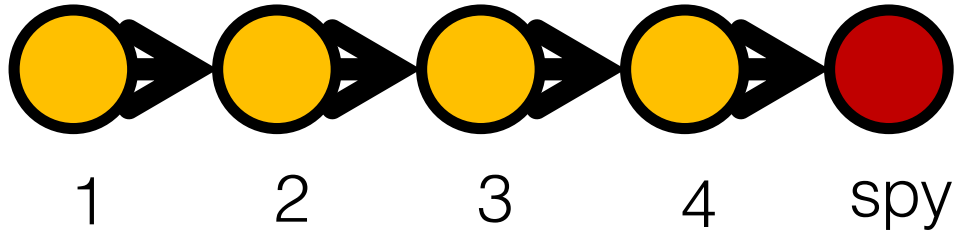

What can we control?

Spreading Protocol

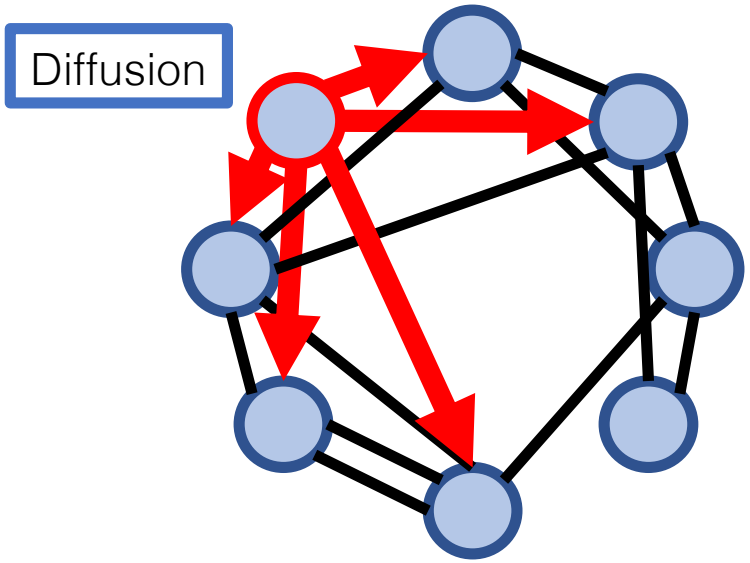

Given a graph, how do we spread content?

Topology

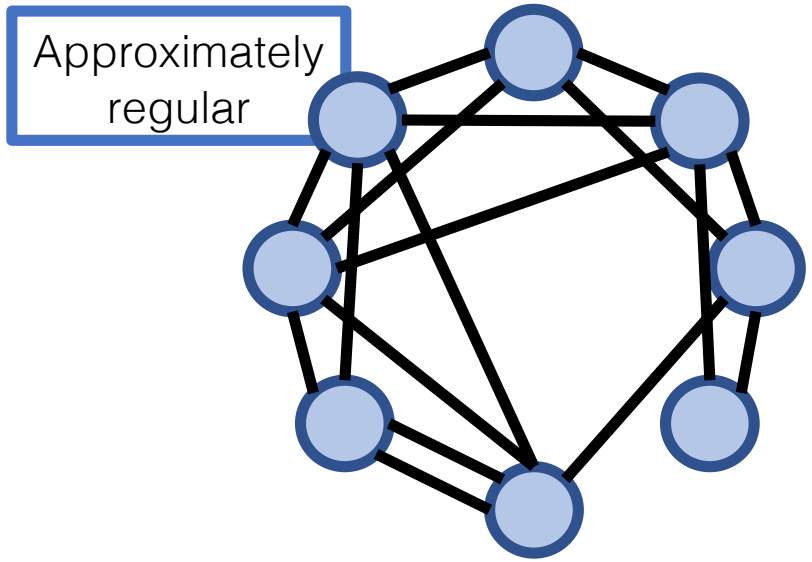

What is the underlying graph topology?

Dynamicity

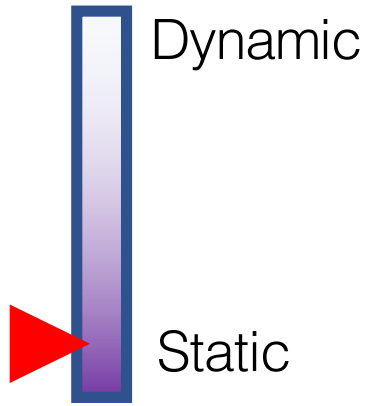

How often does the graph change?

Spreading Protocol: Dandelion

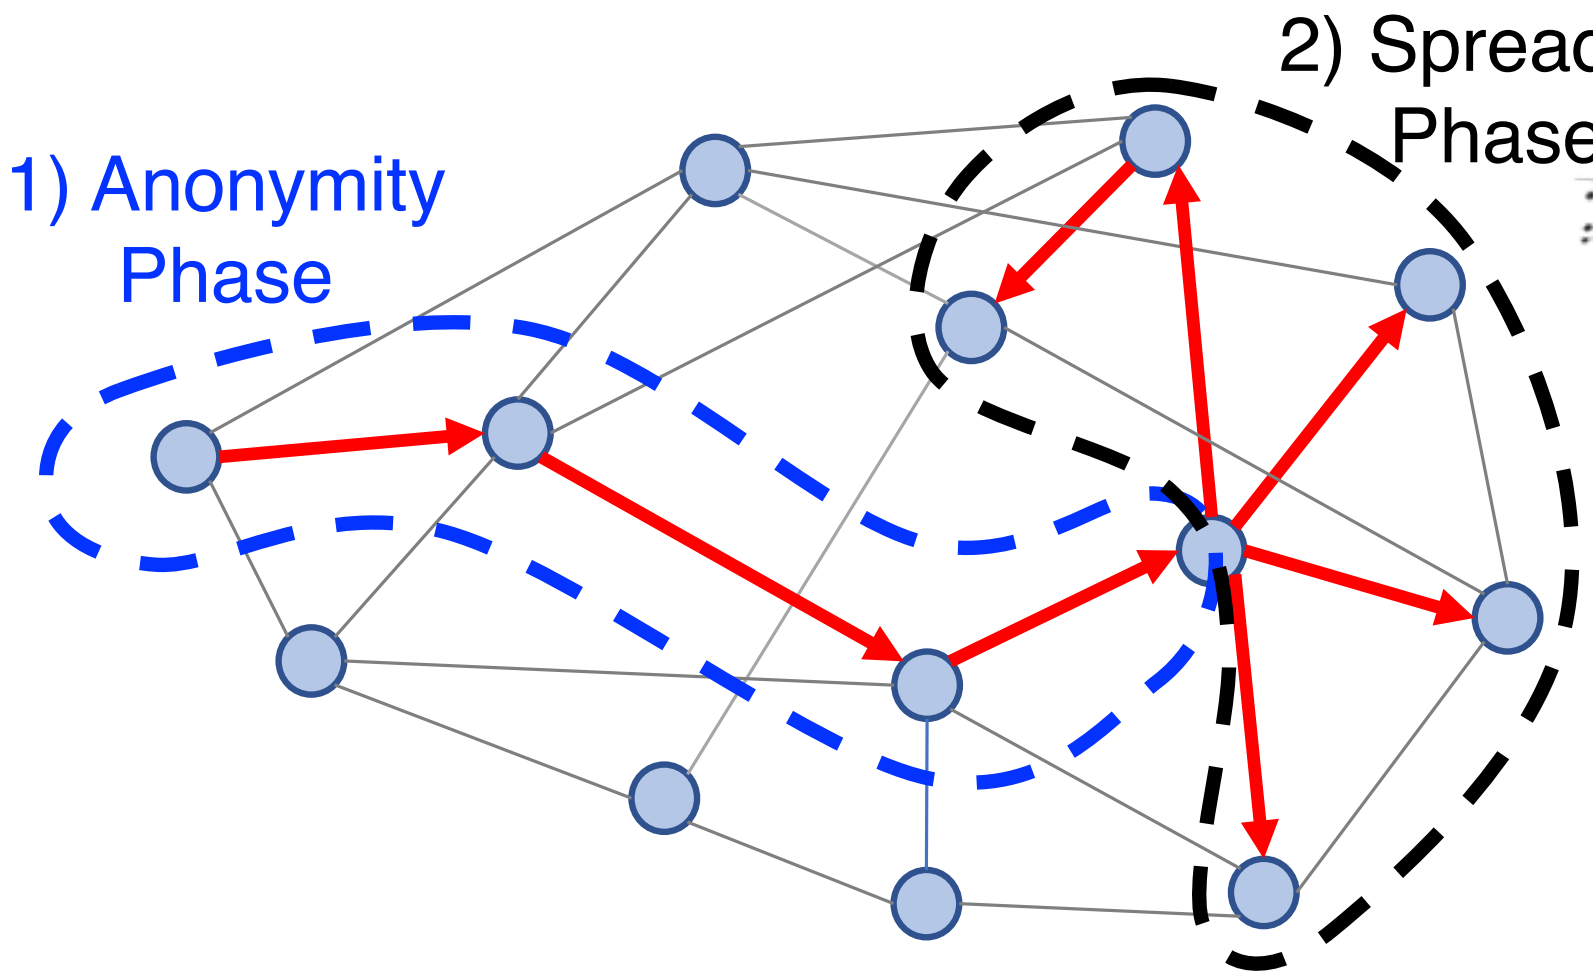

Why Dandelion spreading?

Theorem: Dandelion spreading has an optimally low maximum recall of $p + O\left(\frac{1}{n}\right)$.

lower bound = p

fraction of spies

number of nodes

Graph Topology: Line

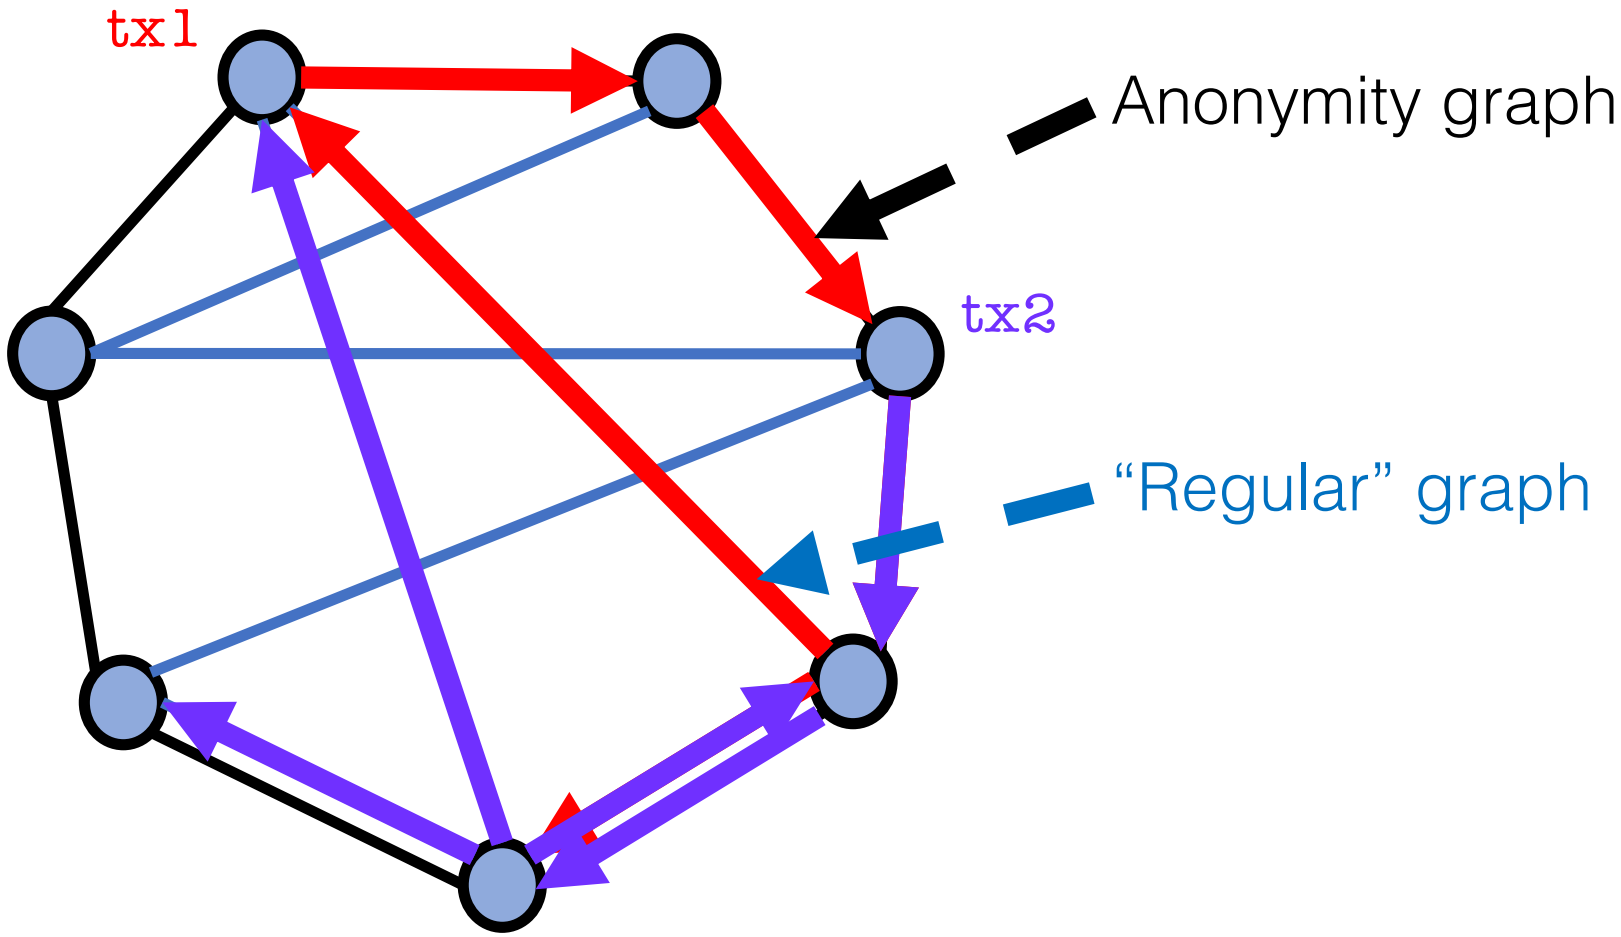

Dynamicity: High

Change the anonymity graph frequently.

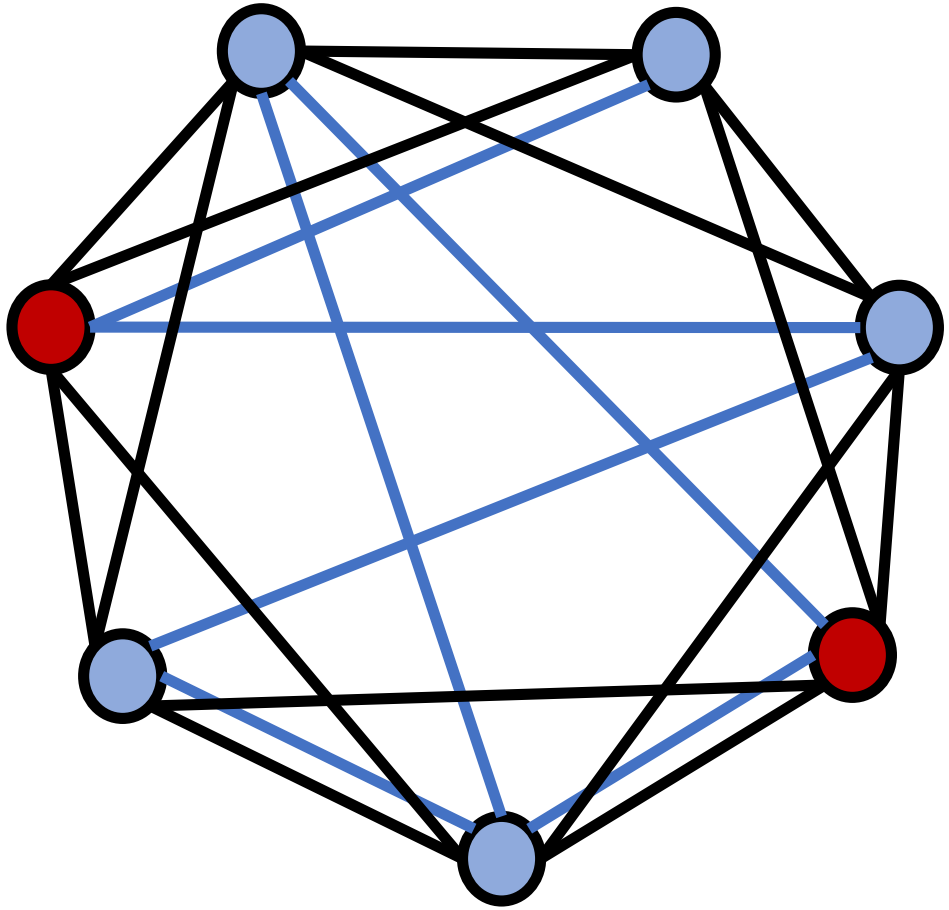

DANDELION Network Policy

Spreading Protocol

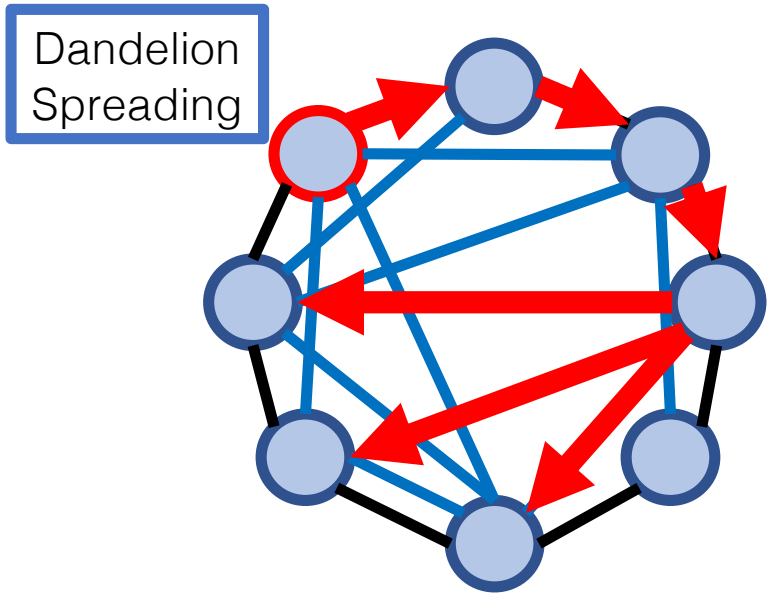

Given a graph, how do we spread content?

Topology

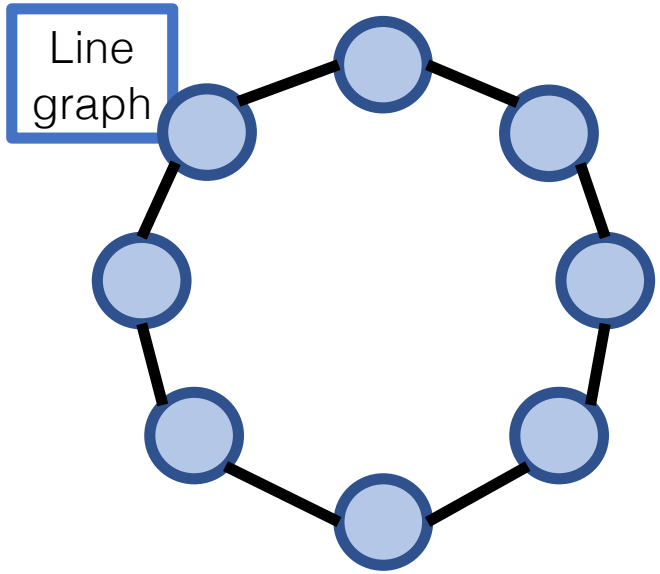

What is the anonymity graph topology?

Dynamicity

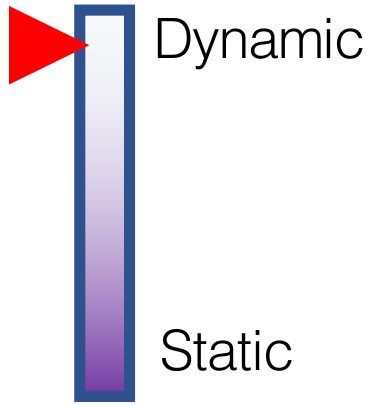

How often does the graph change?

lower bound = p^2

Theorem: DANDELION has a nearly-optimal maximum precision of $\frac{2p^2}{1-p}\log\binom{2}{p} + O\binom{1}{n}$.*

fraction of spies

number of nodes

*For $p < \frac{1}{3}$

Performance: Achievable Region

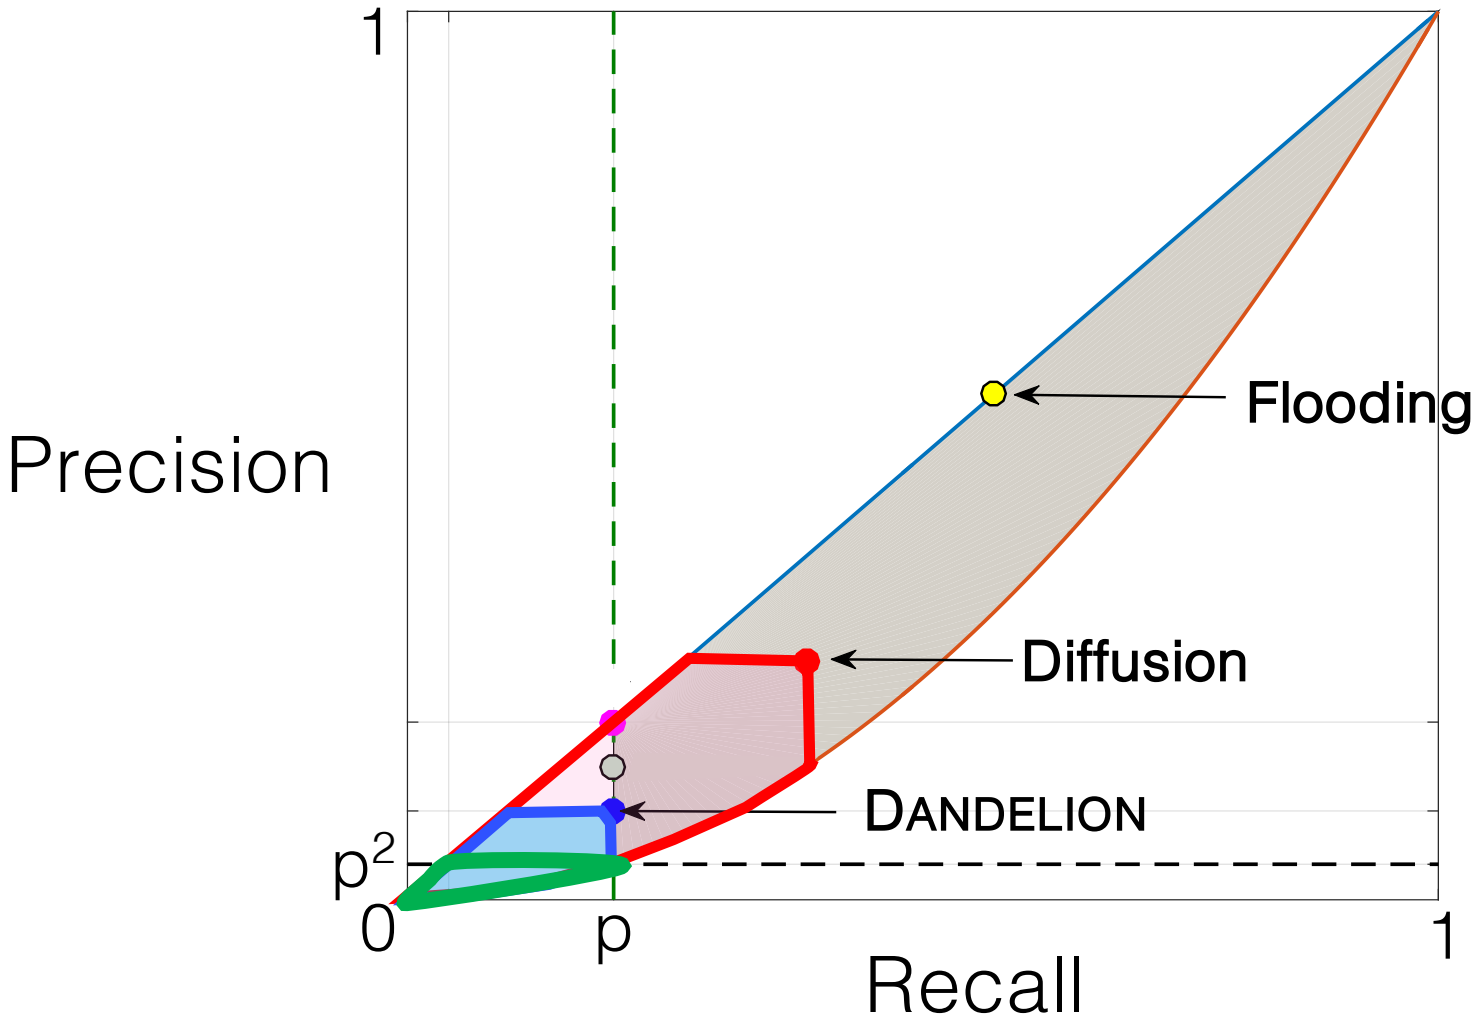

Why does DANDELION work?

Strong mixing properties.

Tree

Precision: $O(p)$

Complete graph
(Crowds, Tor)

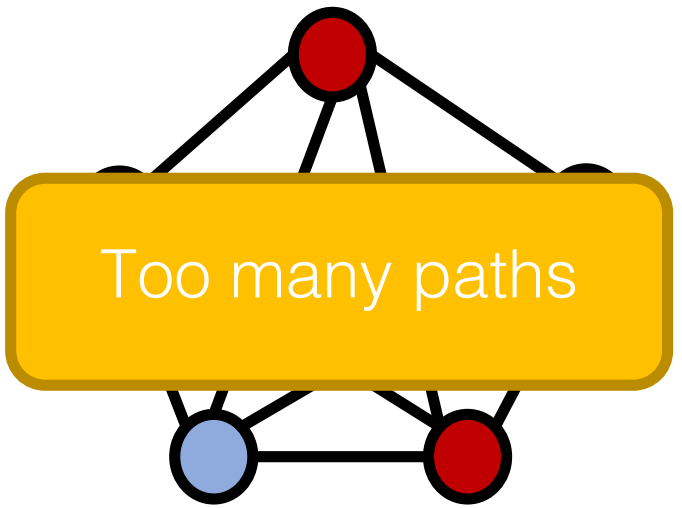

Precision: $\frac{p}{1-p}\,(1-e^{p-1})$

Graph construction in practice

Choose d=1
outbound edges

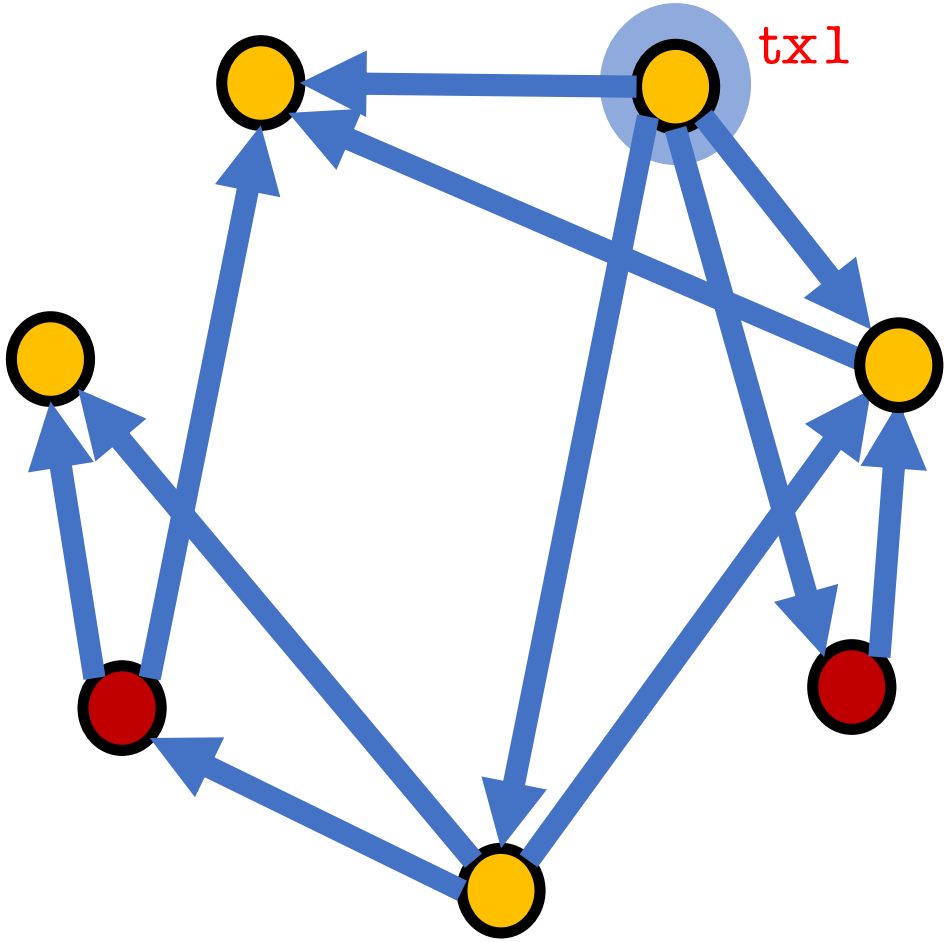

Gives approximate d-regular anonymity graph

d=1

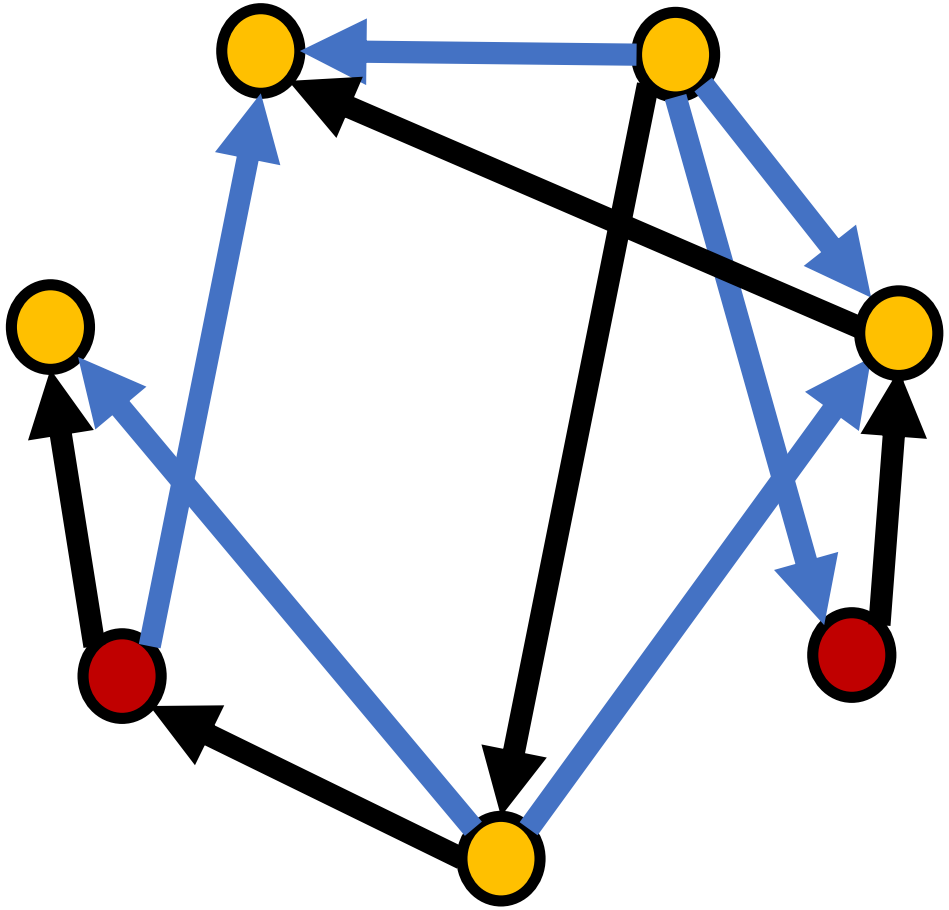

What are drawbacks of Dandelion?

Attack	Effect on Dandelion	Proposed Solution	Effect
Graph Learning	Precision increases to $O(p)$	4-regular anonymity graph	Limits precision gain (Thm. 1)
Intersection	Empirical precision increase	Pseudorandom forwarding	Improved robustness (Thm. 2)
Graph construction	Empirical precision increase	Non-interactive construction	Reduced precision gain
Black hole	Transactions do not propagate	Random stem timers	Provides robustness (Prop. 3)
Partial deployment	Arbitrary recall increase	Blind stem selection	Reduces recall (Thm. 3)

Dandelion++: Lightweight Cryptocurrency Networking with Formal Anonymity Guarantees, ACM Sigmetrics 2018

Experiments on mainnet

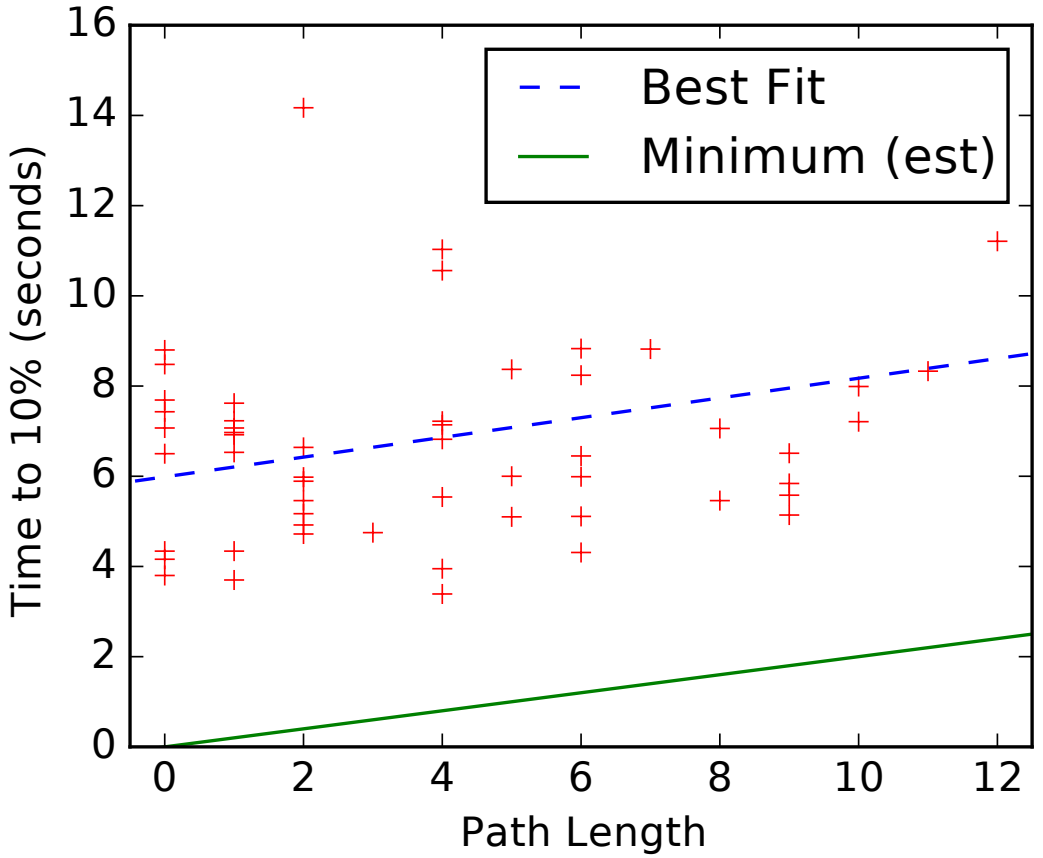

Take-Home Messages

1) Bitcoin's P2P network has weak anonymity protections

2) DANDELION may be a lightweight solution against large-scale deanonymization attacks (but doesn't replace Tor!)

3) More information at:

https://github.com/dandelion-org/bips
https://github.com/dandelion-org/bitcoin

Simulation on Bitcoin P2P Topology

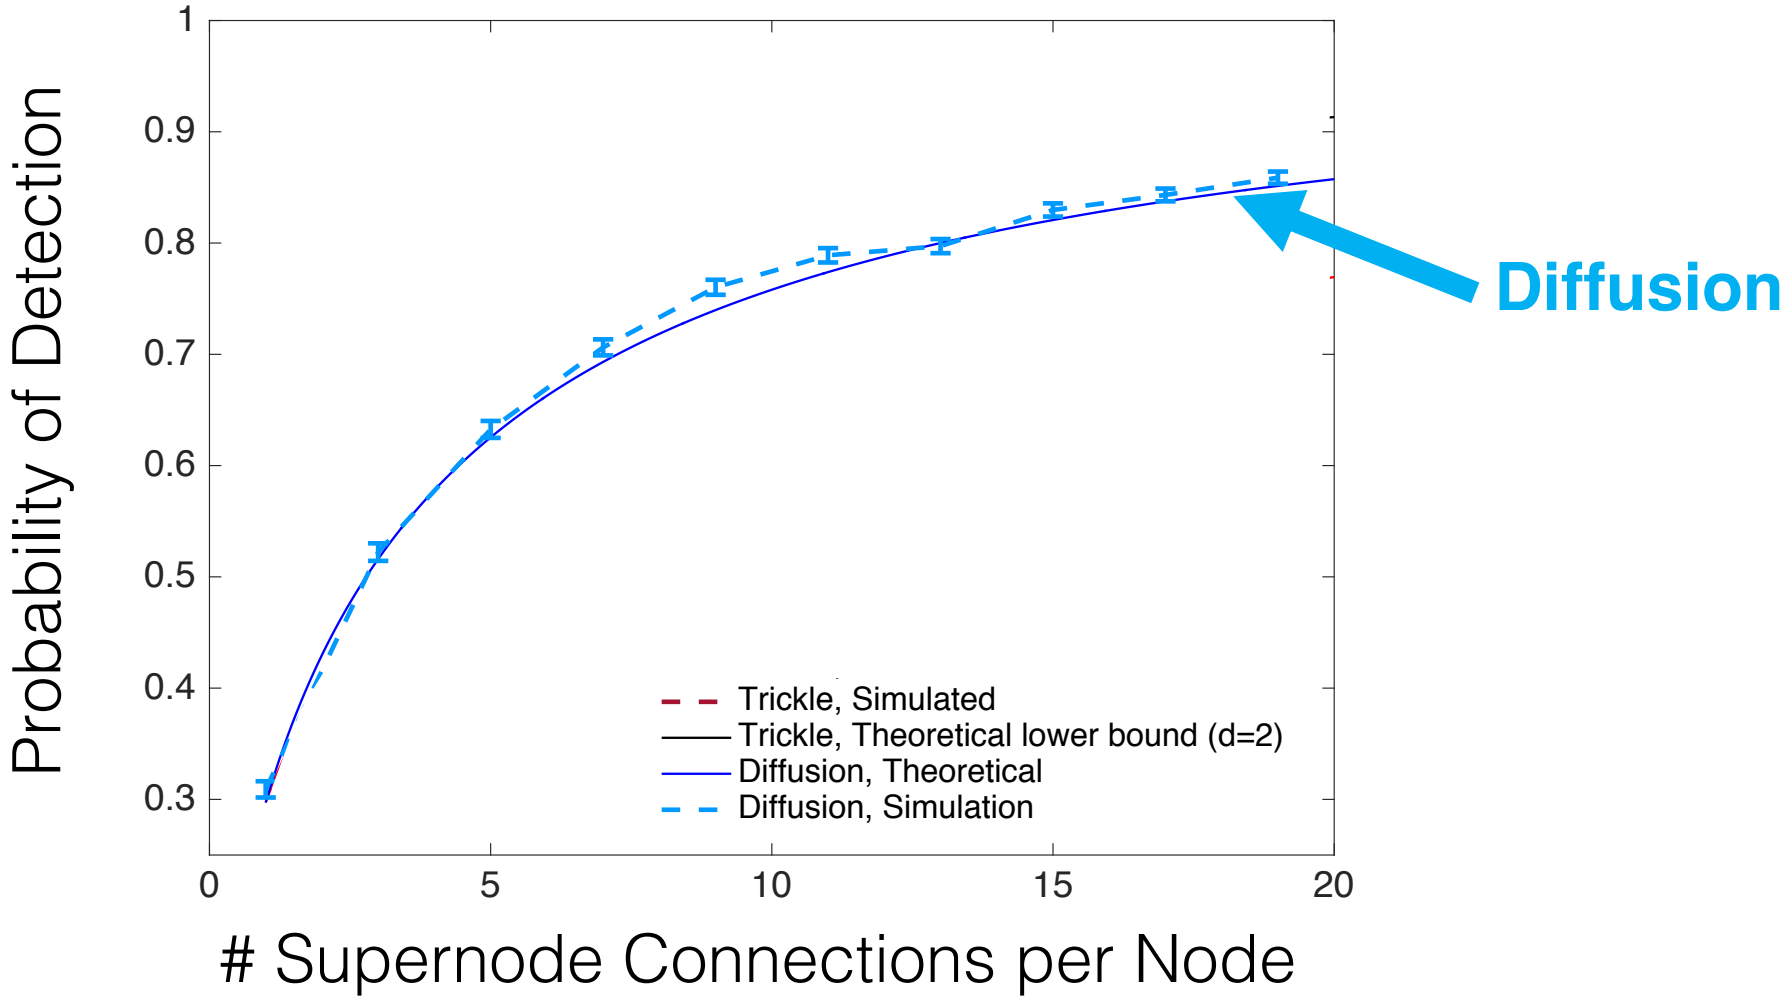

F. and Viswanath, NIPS 2017

4-Regular Graphs

- More robust against adversaries that learn the graph

- Per-transaction routing vulnerable to intersection attacks

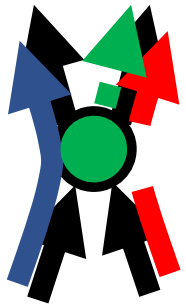

One-to-one
Routing

- **Pro**: Increases cost of graph-learning attacks
- **Con**: Can make transactions from the same source easier to link

FAQ: Why not Tor?

- Tor, VPNs, etc. address this problem
- Only work for savvy or privacy-aware users
- If Bitcoin is to become a mainstream payment system, it should protect everyone's transactions
- Dandelion: lightweight, easy to integrate into existing network

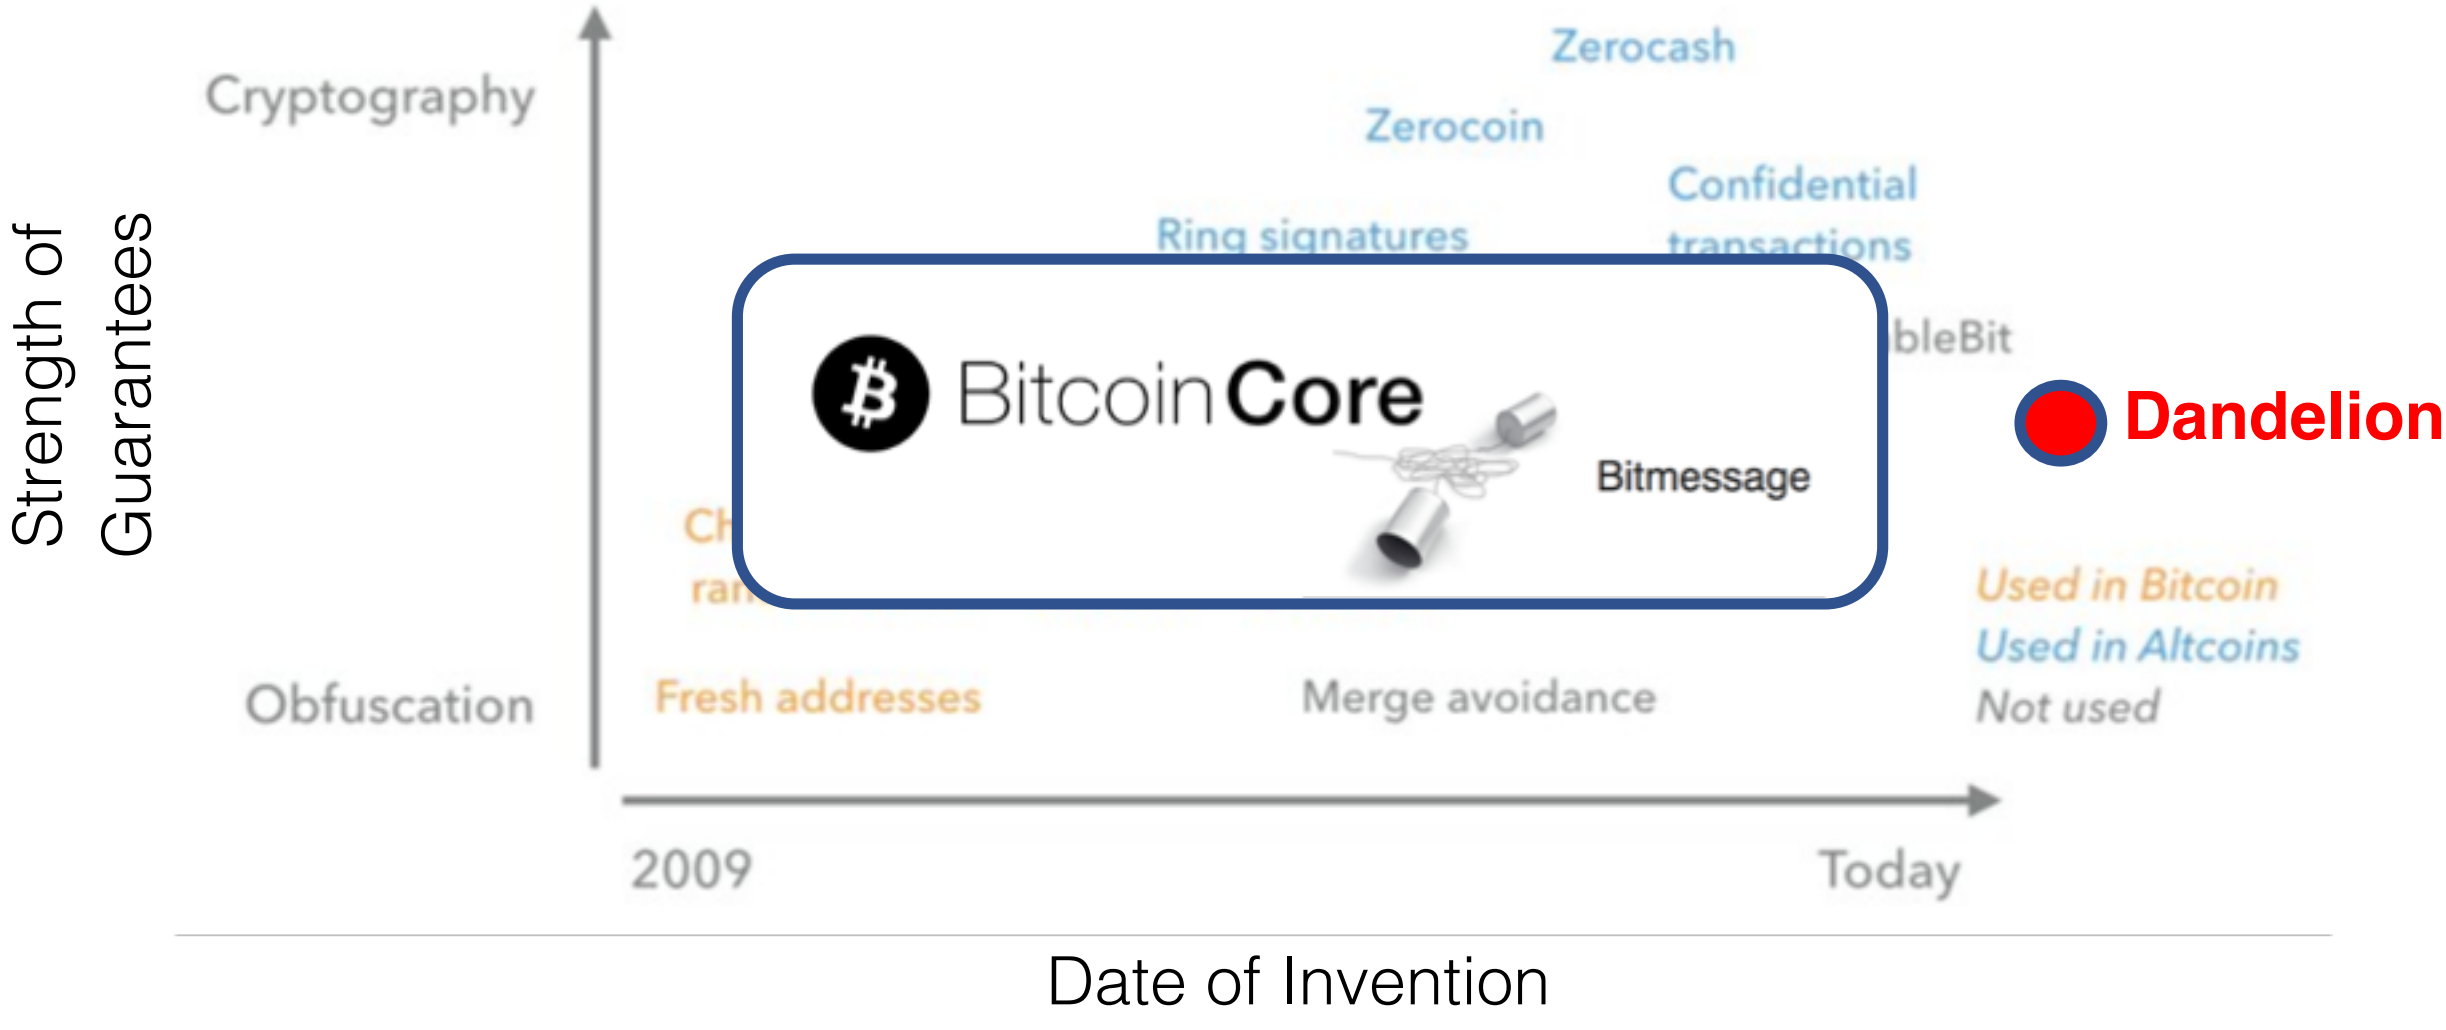

Narayanan and Möser, 2017

Moving from theory to practice

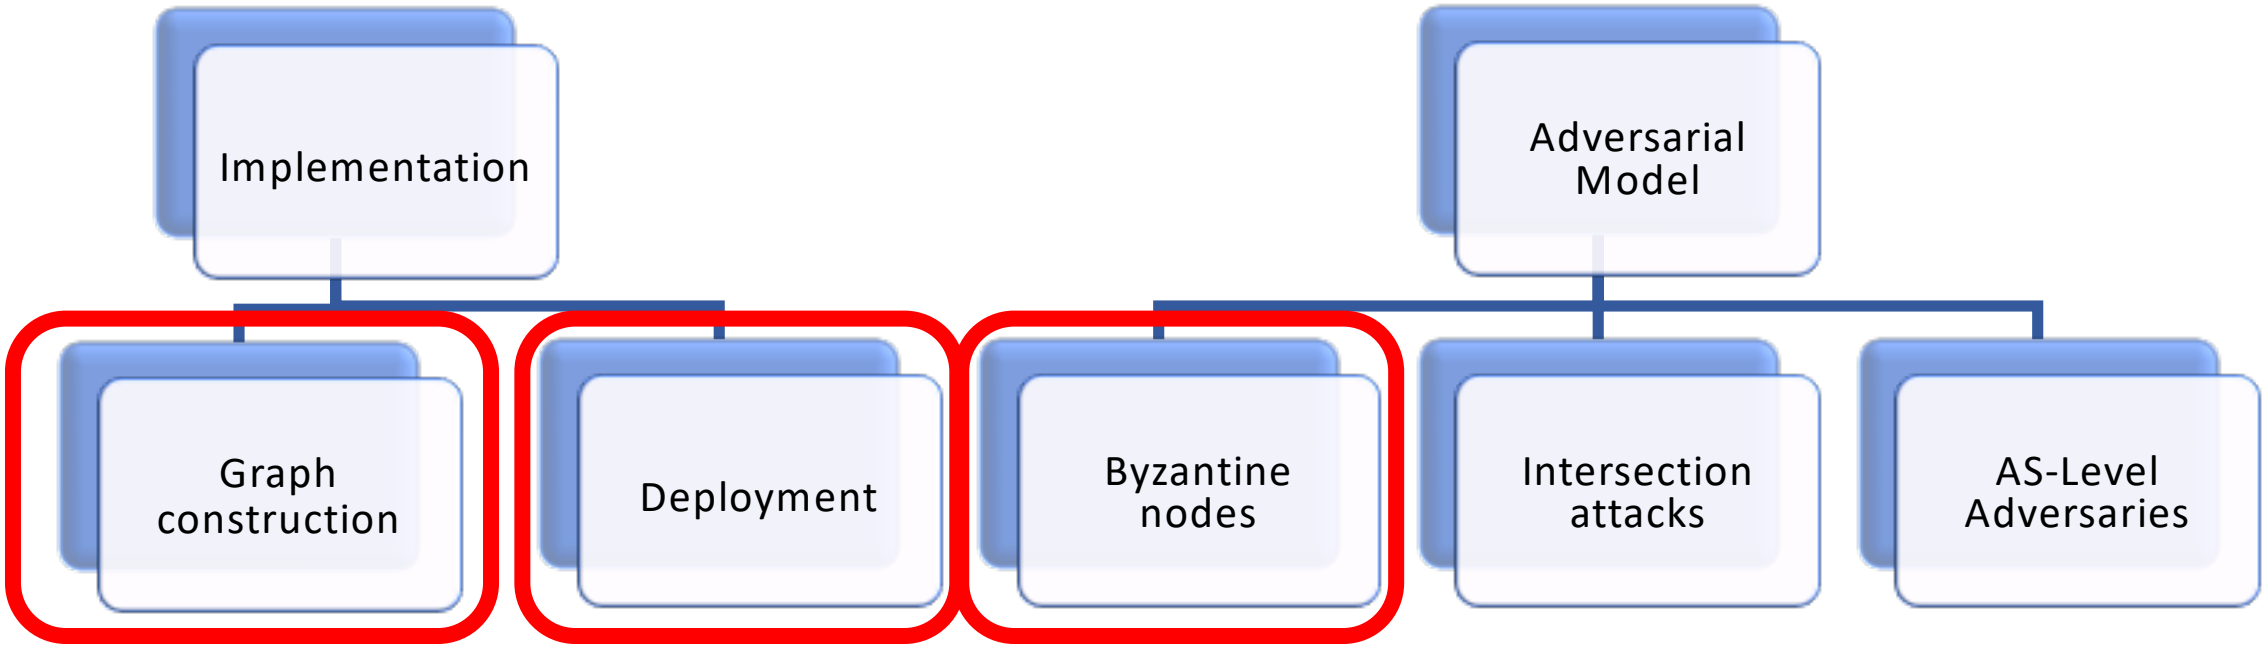
Implementation
Graph construction
Deployment
Adversarial Model
Byzantine nodes
Intersection attacks
AS-Level Adversaries

Implementation: Dandelion spreading

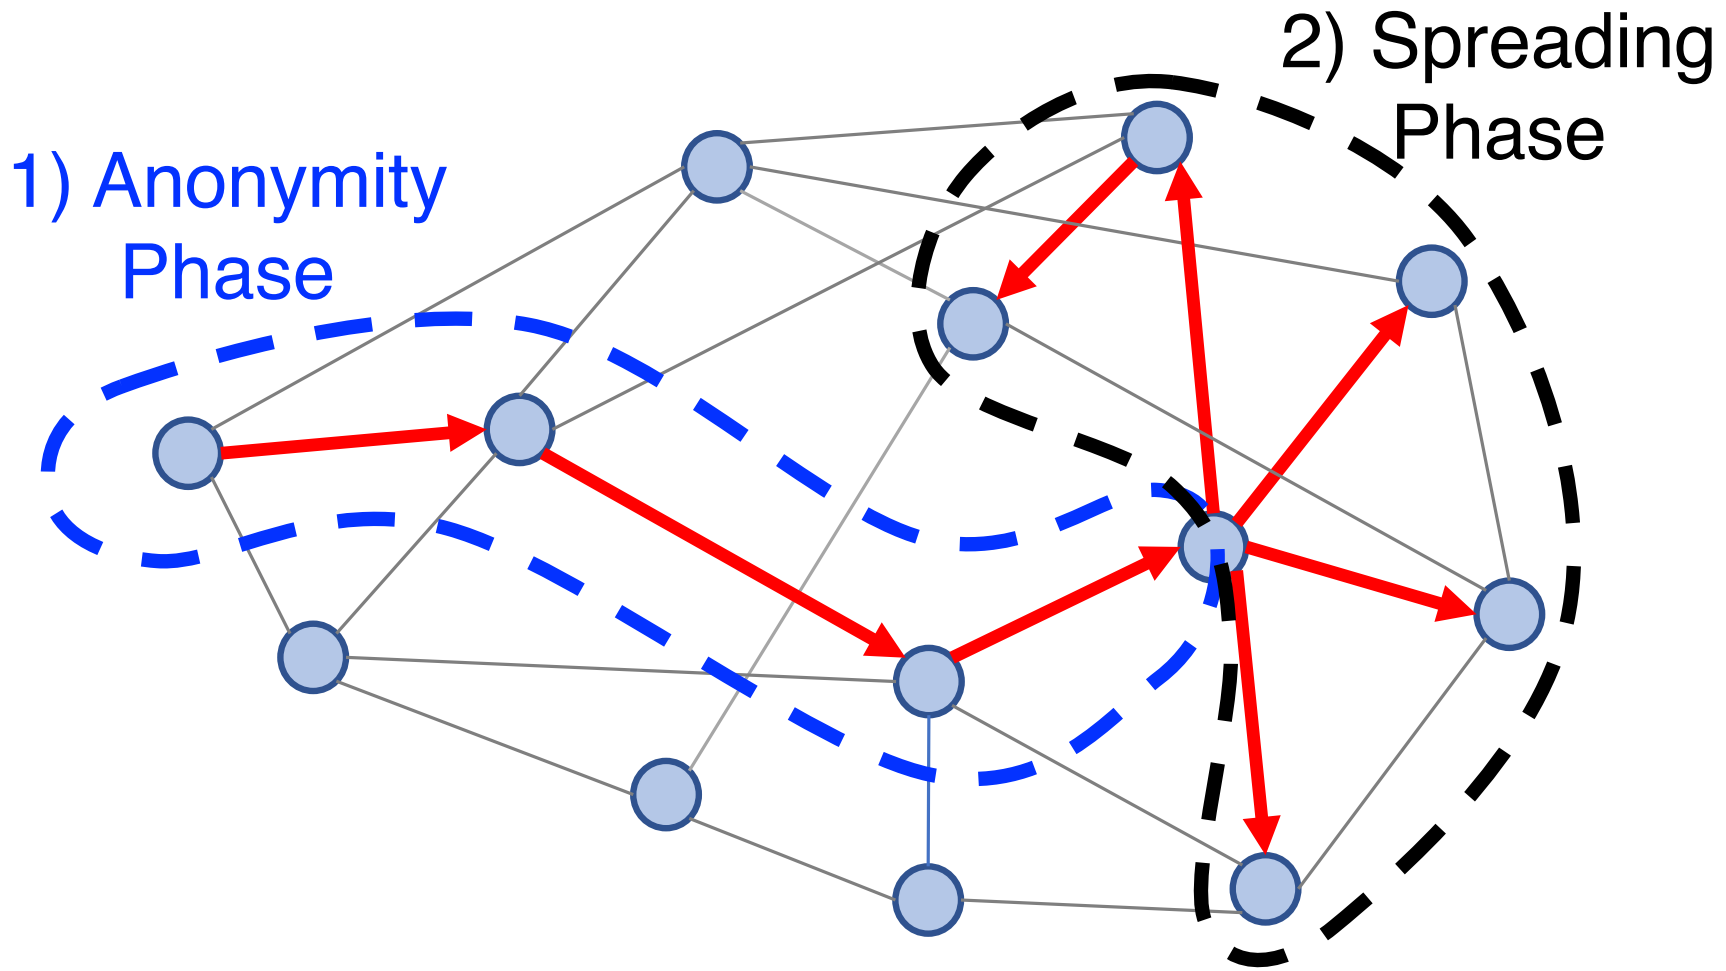

Anonymity graph construction

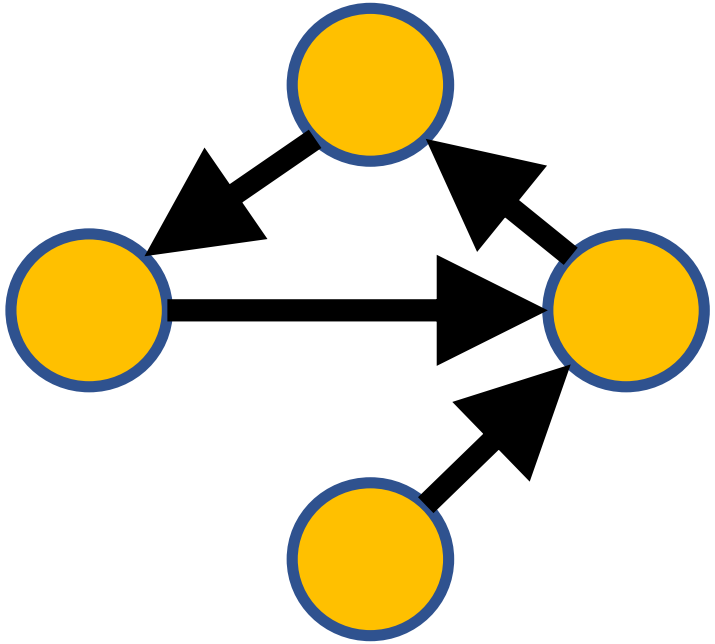

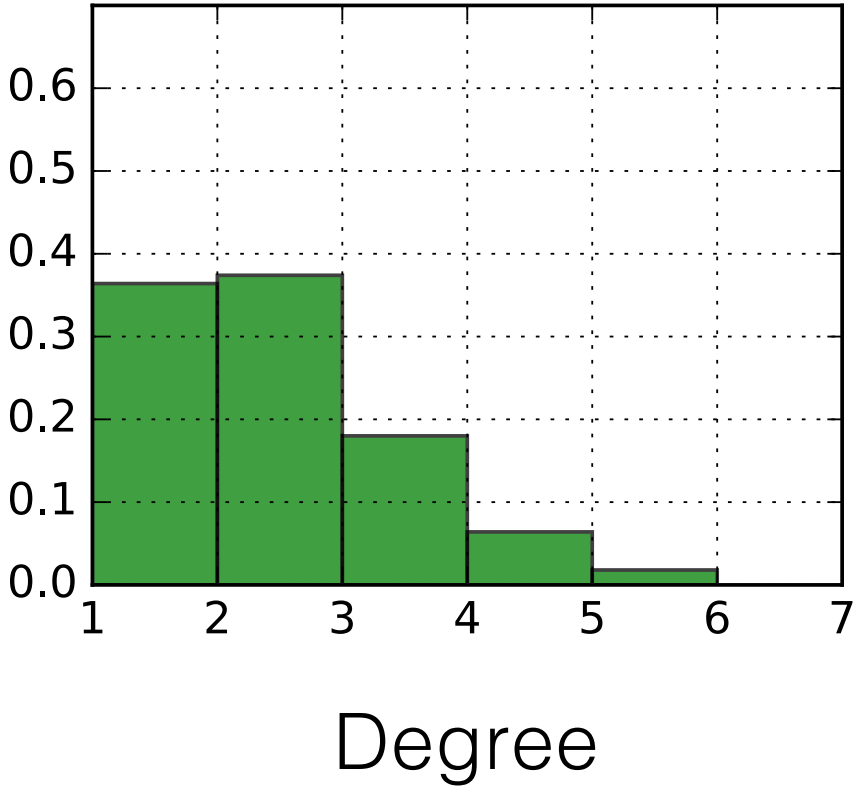

Adversarial Model: Byzantine nodes

Learn the graph

4-regular graphs

Misbehave during graph construction

Misbehave during propagation

Anonymity graph construction

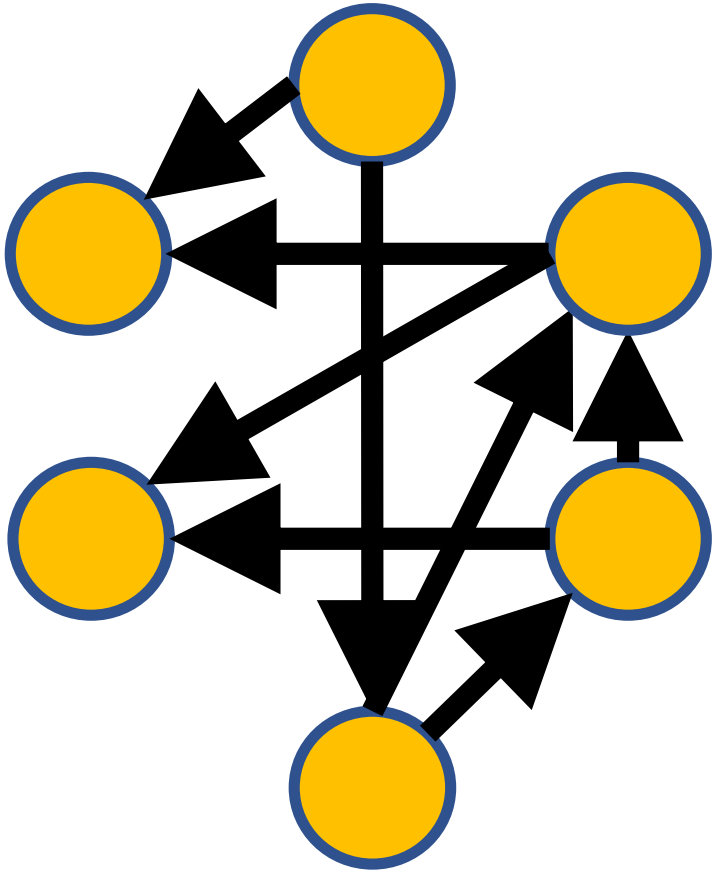

Dealing with stronger adversaries

Learn the graph

4-regular graphs

Misbehave during graph construction

Only send messages on outgoing edges

Misbehave during propagation

Multiple nodes diffuse

Partial deployment

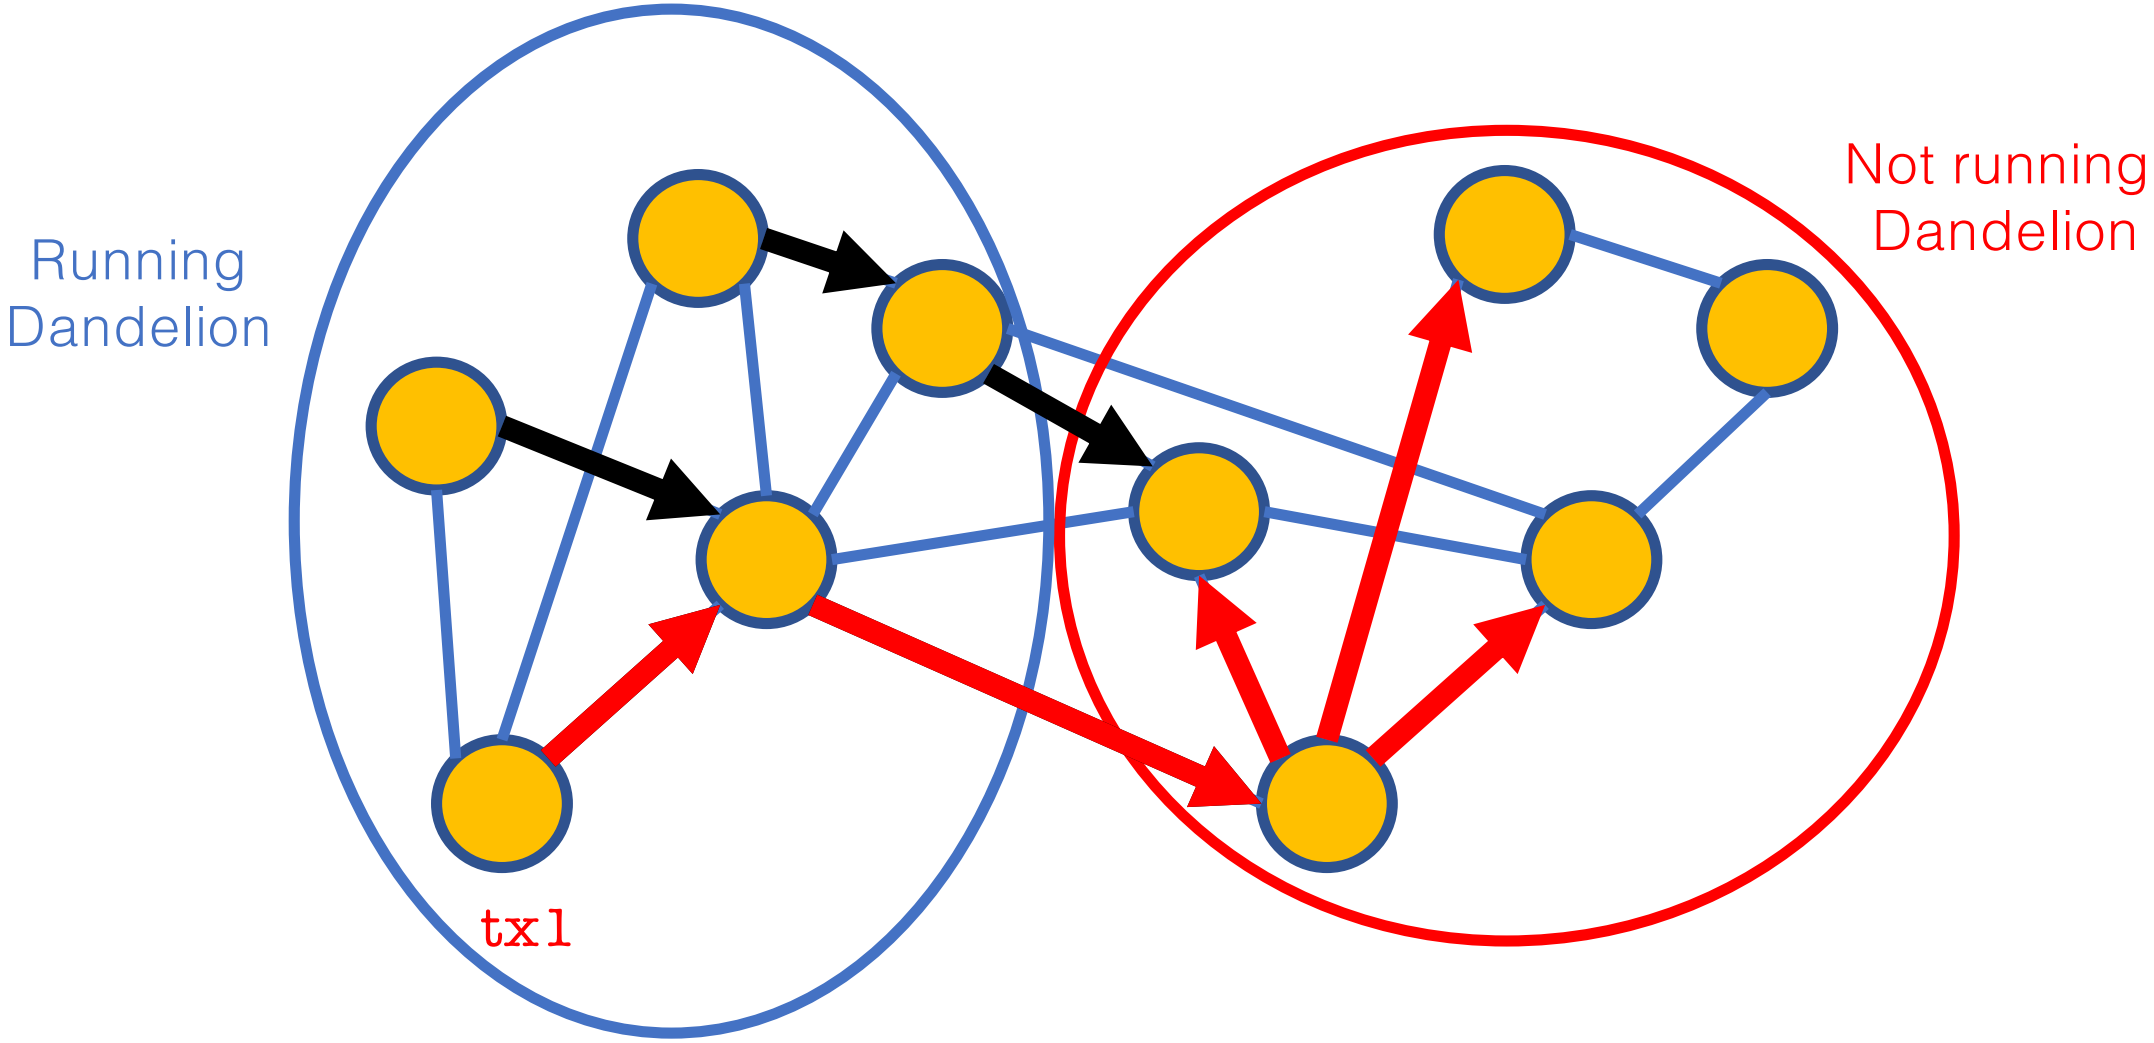

Latency Overhead: Estimate

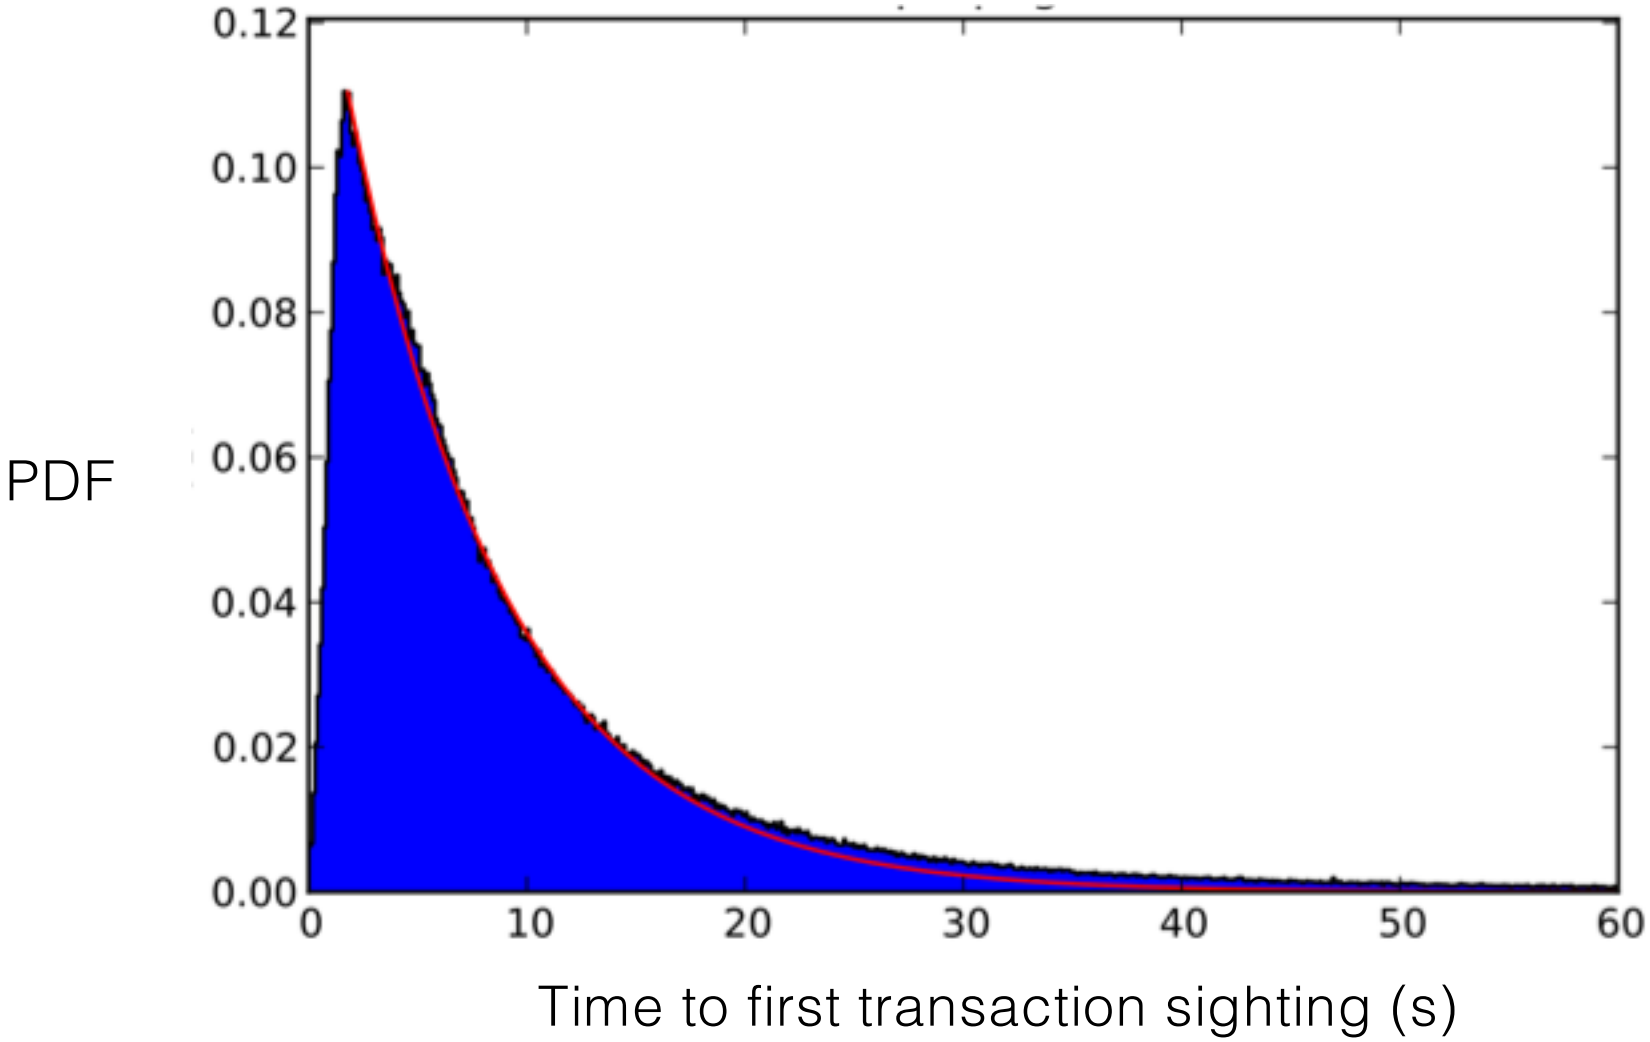

Information Propagation in the Bitcoin Network, Decker and Wattenhofer, 2013

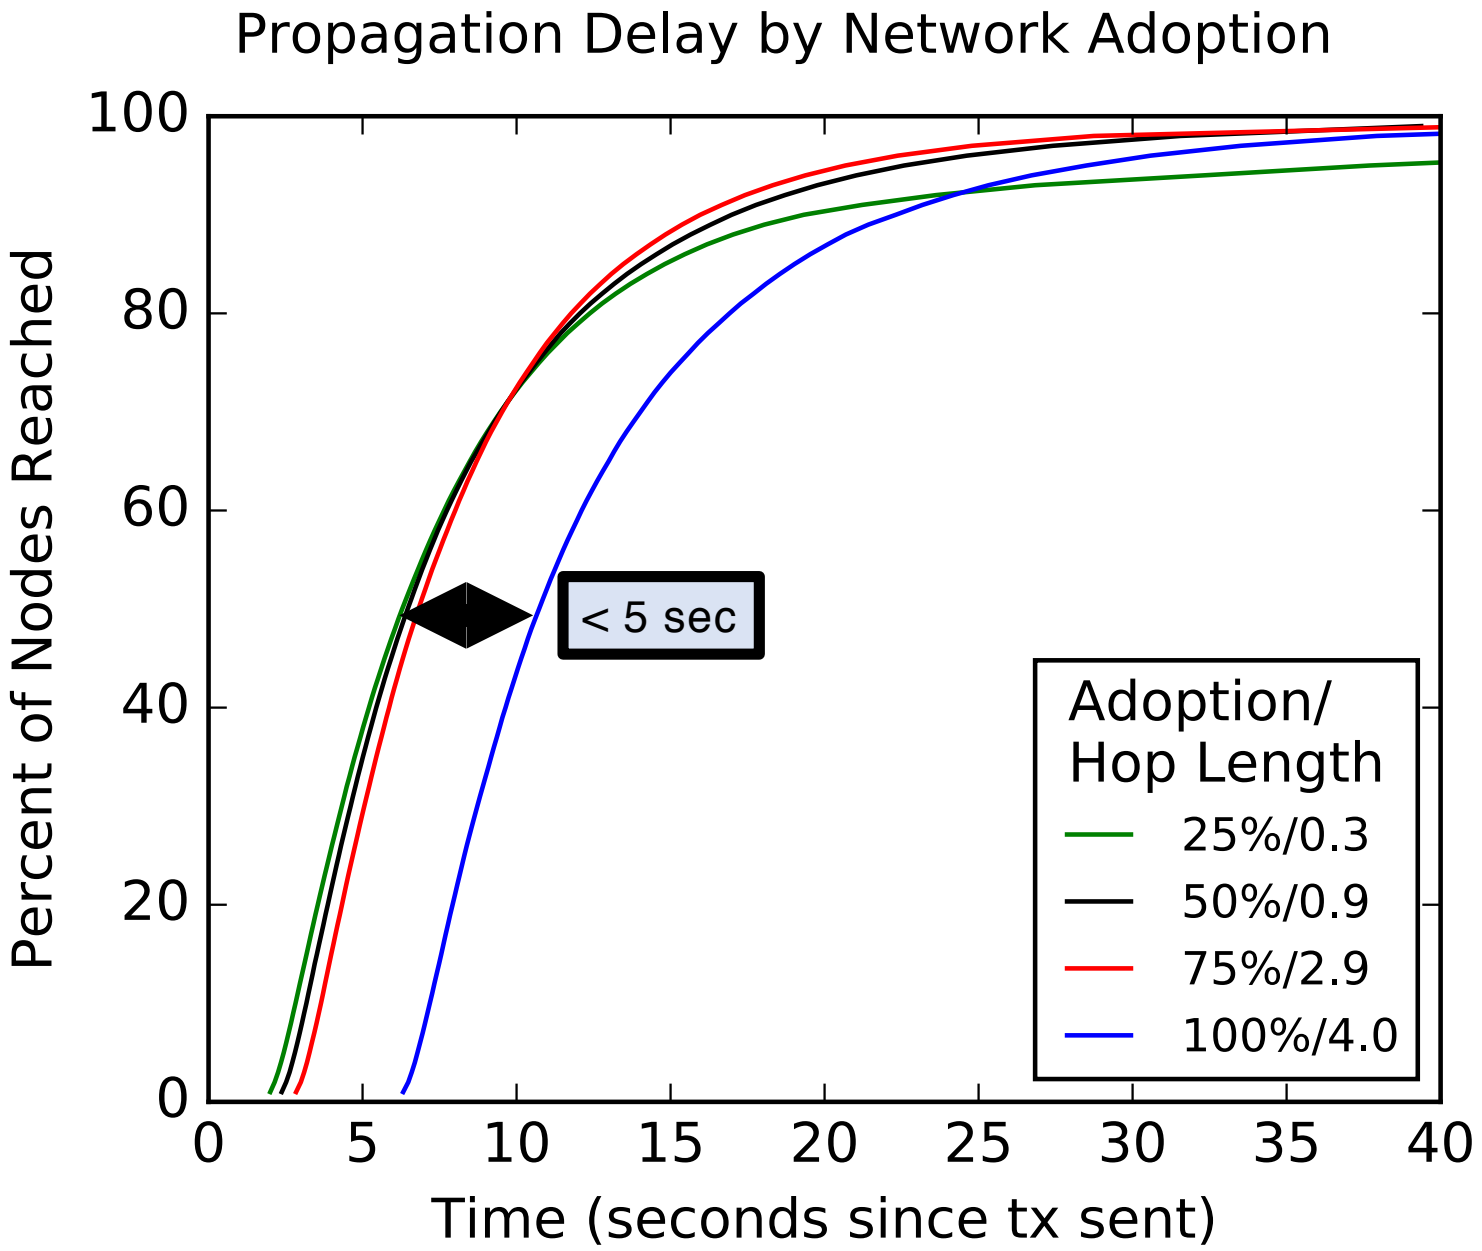
Propagation Delay by Network Adoption
Percent of Nodes Reached
Time (seconds since tx sent)
100
80
60
40
20
0
0
5
10
15
20
25
30
35
40
< 5 sec
Adoption/
Hop Length
25%/0.3
50%/0.9
75%/2.9
100%/4.0

DANDELION vs. Tor, Crowds, etc.

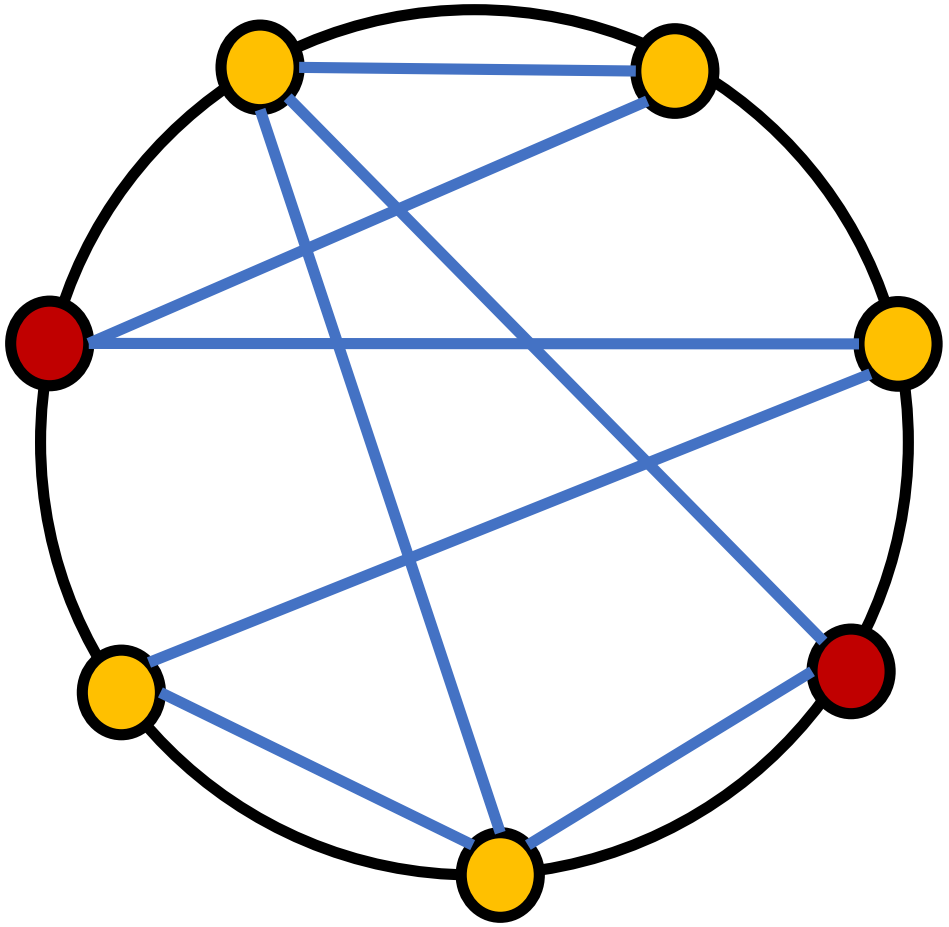

1) Messages propagate over the **same** cycle graph

2) Anonymity graph changes dynamically.

3) No encryption required.

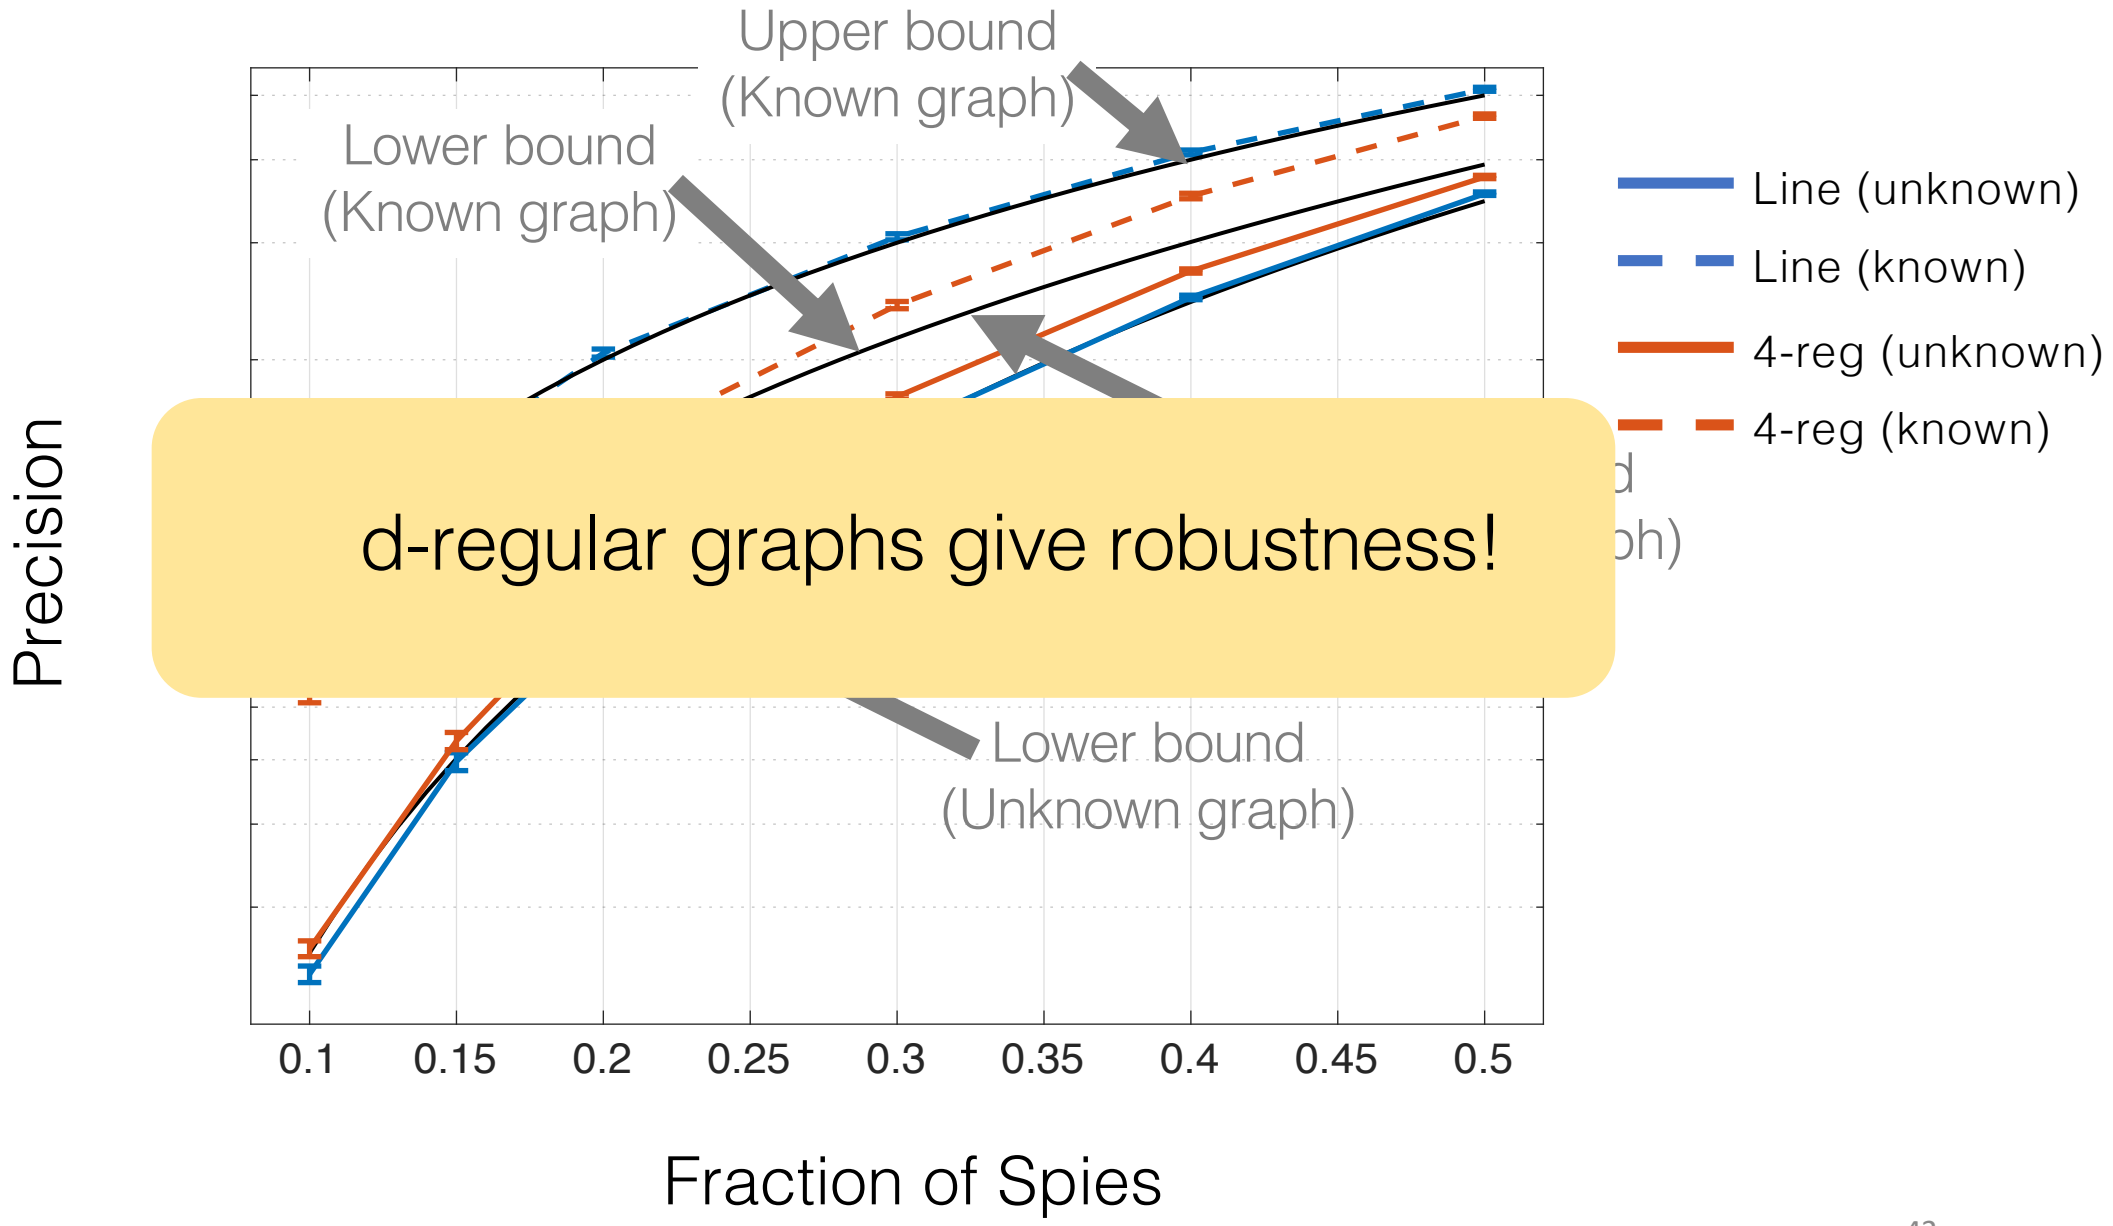

d-regular graphs give robustness!

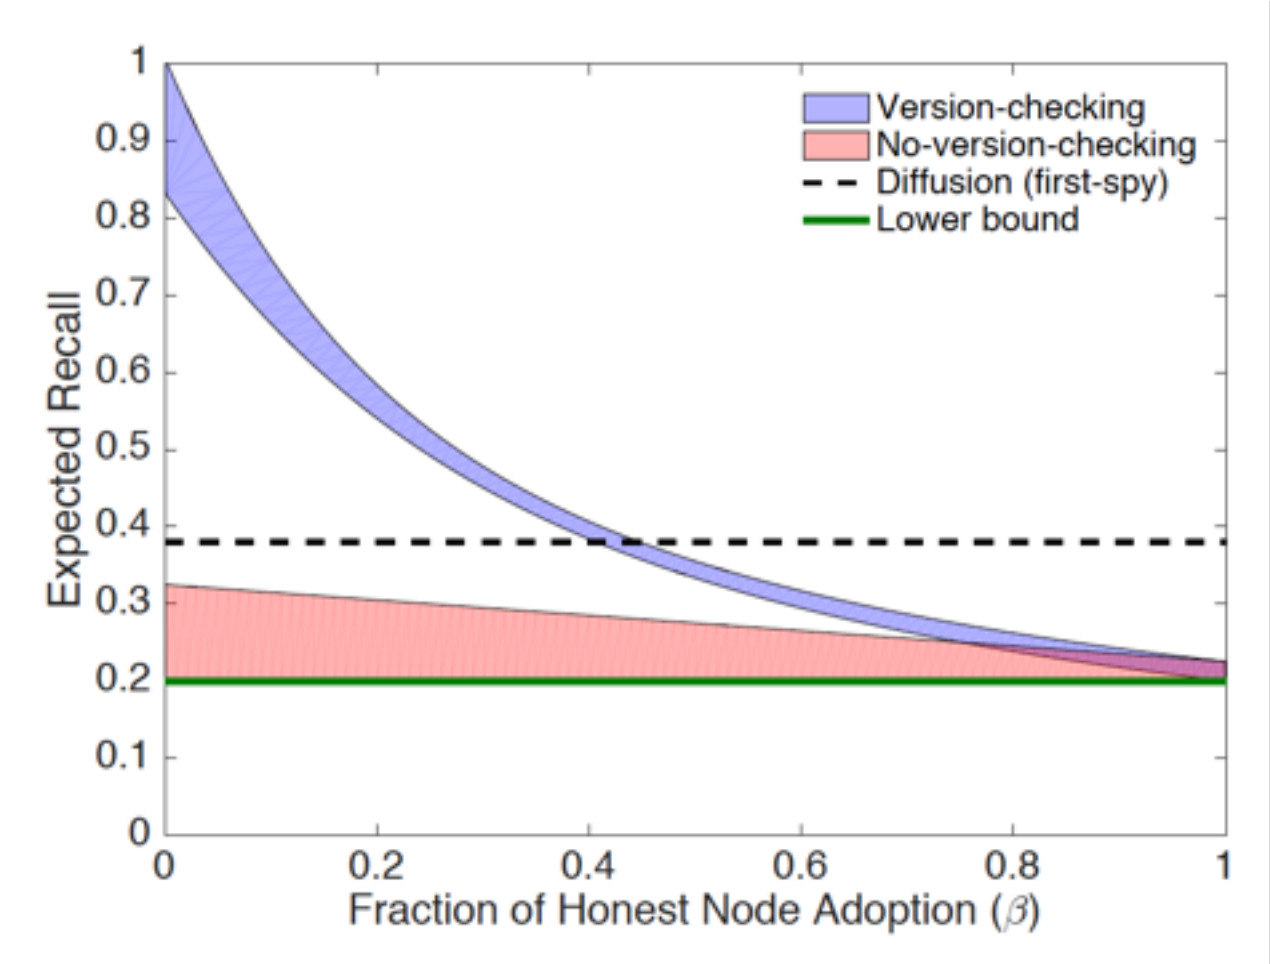
Version-checking
No-version-checking
Diffusion (first-spy)
Lower bound
Expected Recall
Fraction of Honest Node Adoption (β)
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
0
0.2
0.4
0.6
0.8
1

Anonymity graph construction

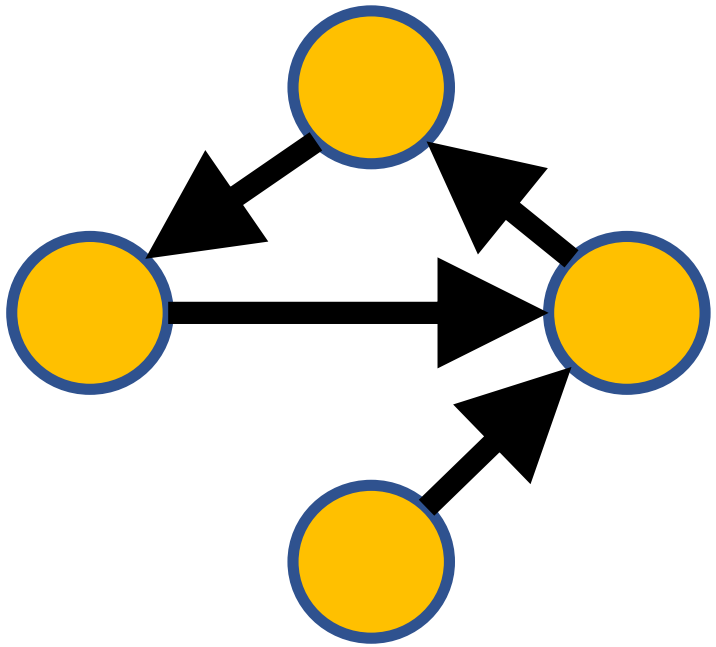

Base Case

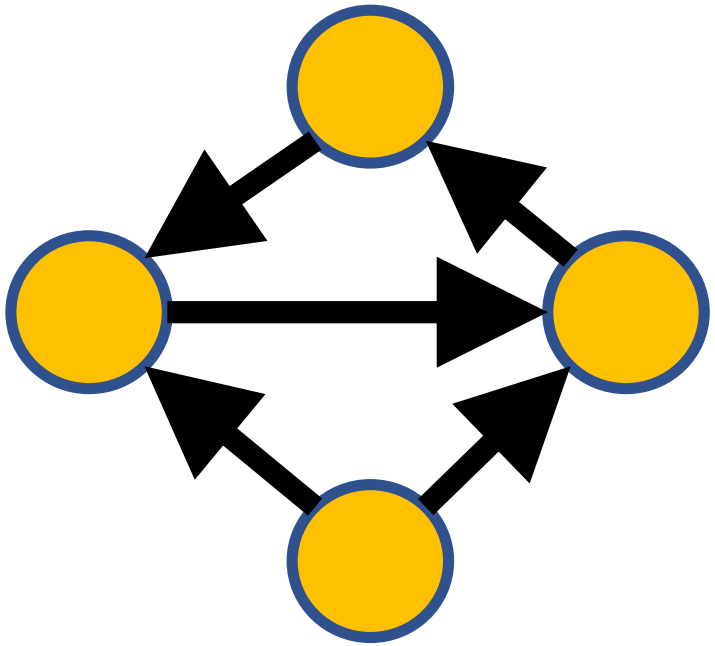

k=1 rounds of Degree-Checking

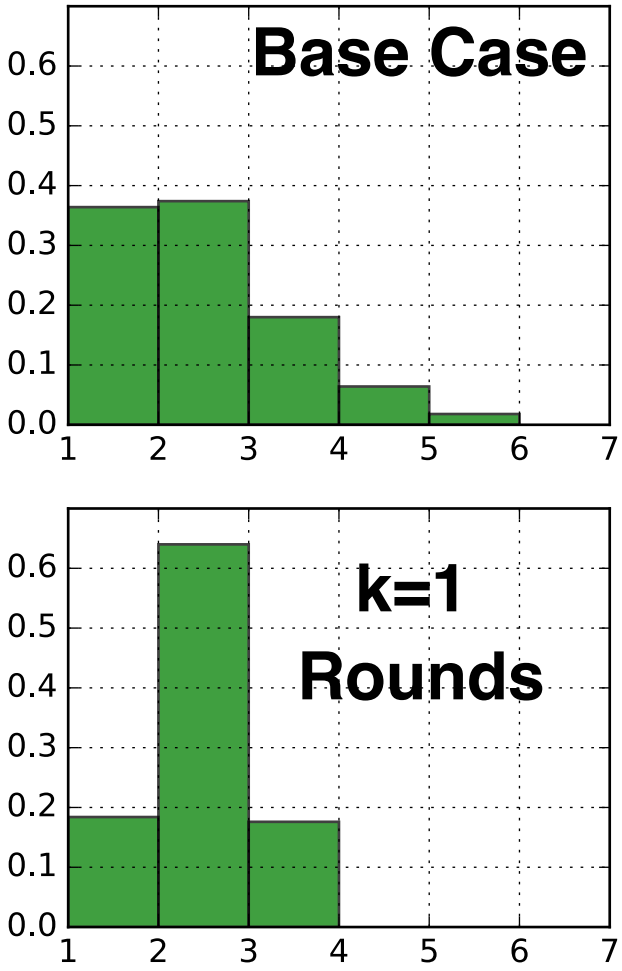

Degree

Dealing with stronger adversaries

Learn the graph	Misbehave during graph construction	Misbehave during propagation
⬇	⬇	⬇
4-regular graphs	Get rid of degree-checking	Multiple nodes diffuse

Learning the anonymity graph

Precision

	Line	Random regular
	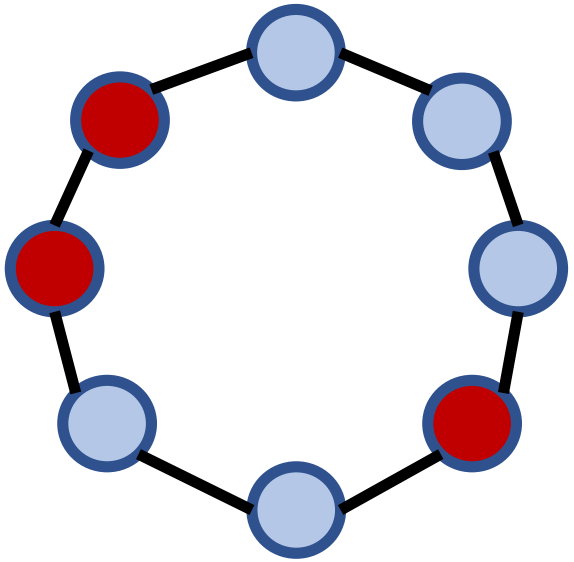	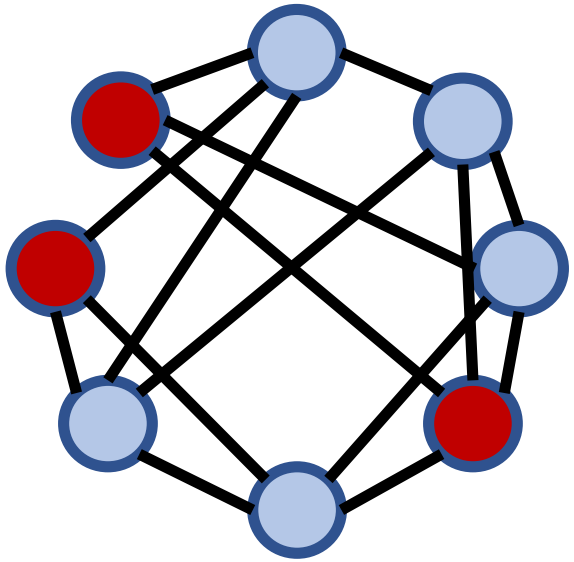
Graph unknown	$O\left(\mathrm{p}^2 \log\left(\frac{1}{p}\right)\right)$	?
Graph known	$\Omega(p)$	

Manipulating the anonymity graph

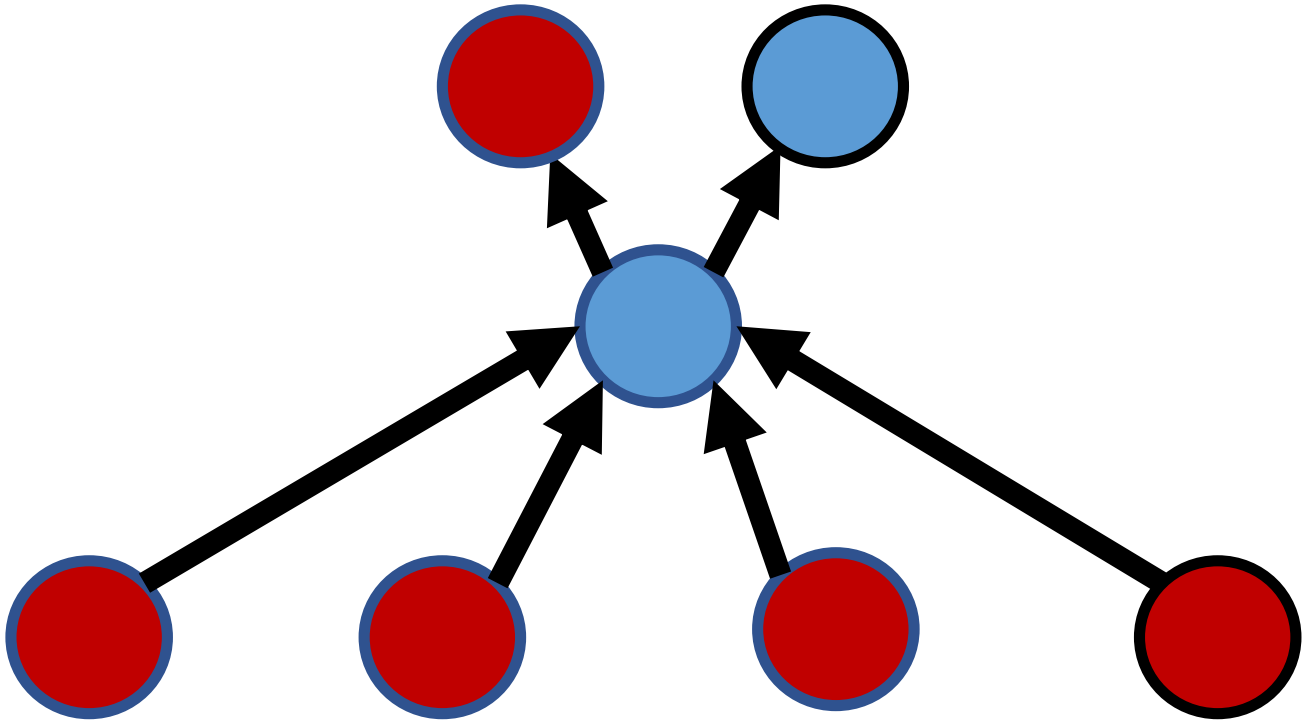

DANDELION++ Network Policy

Spreading Protocol

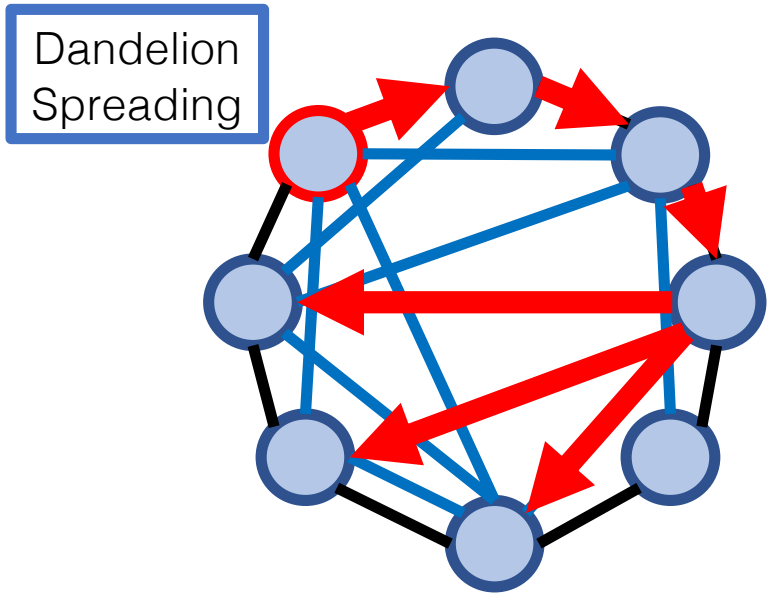

Given a graph, how do we spread content?

Topology

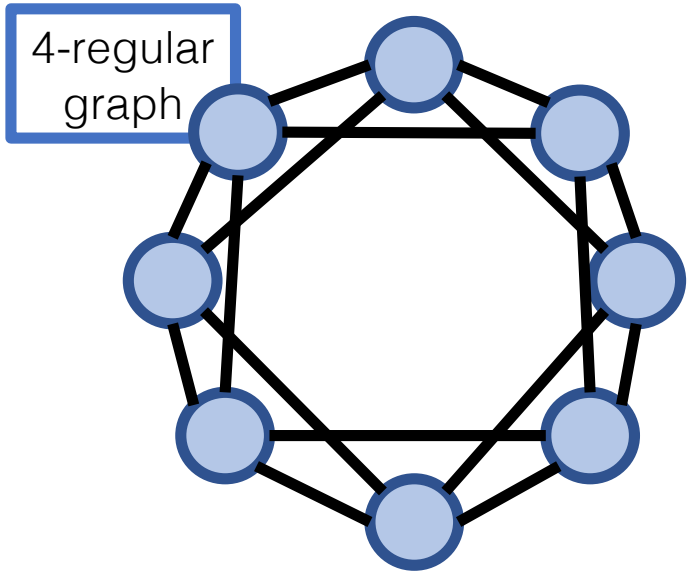

What is the anonymity graph topology?

Dynamicity

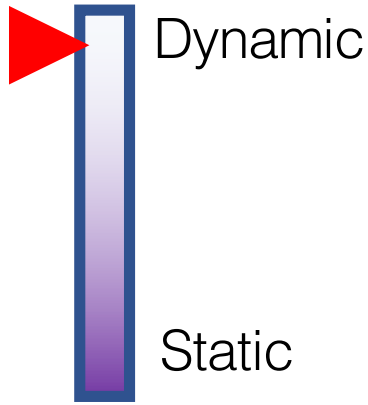

How often does the graph change?